Matrix DFE Series and N-SA User's Guide

Important Notice

This guide is a work-in-progress. It is being made available to provide information about key configuration tasks, but it does not yet represent the full functionality of Matrix DFE Series and N-SA devices. Updated versions of this guide with additional chapters will be posted on the Enterasys Networks website as they become available. Please refer to the website and the revision history table below to determine if a newer version has been published.

Revision History

Book Version Number	Date	Description of Changes
9033939	July 2004	Original document.
9033939-01	February 2005	Added 5.01 firmware changes, style changes and new chapters.
9033939-02	February 2006	Removed erroneous information about SSHv2 and Telnet from Chapter 3.

9033939-02

enterasys™
Networks that Know

Notice

Enterasys Networks reserves the right to make changes in specifications and other information contained in this document and its web site without prior notice. The reader should in all cases consult Enterasys Networks to determine whether any such changes have been made.

The hardware, firmware, or software described in this document is subject to change without notice.

IN NO EVENT SHALL ENTERASYS NETWORKS BE LIABLE FOR ANY INCIDENTAL, INDIRECT, SPECIAL, OR CONSEQUENTIAL DAMAGES WHATSOEVER (INCLUDING BUT NOT LIMITED TO LOST PROFITS) ARISING OUT OF OR RELATED TO THIS DOCUMENT, WEB SITE, OR THE INFORMATION CONTAINED IN THEM, EVEN IF ENTERASYS NETWORKS HAS BEEN ADVISED OF, KNEW OF, OR SHOULD HAVE KNOWN OF, THE POSSIBILITY OF SUCH DAMAGES.

Enterasys Networks, Inc.
50 Minuteman Road
Andover, MA 01810

© 2006 Enterasys Networks, Inc. All rights reserved.

Part Number: 9033939-02 February 2006

ENTERASYS, ENTERASYS NETWORKS, ENTERASYS MATRIX, and any logos associated therewith, are trademarks or registered trademarks of Enterasys Networks, Inc. in the United States and other countries. All other product names mentioned in this manual may be trademarks or registered trademarks of their respective owners.

Documentation URL: http://www.enterasys.com/support/manuals

Code Copyrights and Acknowledgements

Copyright (c) 1999-2001 Internet Software Consortium.
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of The Internet Software Consortium nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE INTERNET SOFTWARE CONSORTIUM AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE INTERNET SOFTWARE CONSORTIUM OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software has been written for the Internet Software Consortium by Ted Lemon in cooperation with Vixie Enterprises and Nominum, Inc.

To learn more about the Internet Software Consortium, see “http://www.isc.org/”. To learn more about Vixie Enterprises, see “http://www.vix.com”. To learn more about Nominum, Inc., see “http://www.nominum.com”.

Enterasys Networks, Inc.
Firmware License Agreement

BEFORE OPENING OR UTILIZING THE ENCLOSED PRODUCT,
CAREFULLY READ THIS LICENSE AGREEMENT.

This document is an agreement ("Agreement") between the end user ("You") and Enterasys Networks, Inc. on behalf of itself and its Affiliates (as hereinafter defined) ("Enterasys") that sets forth Your rights and obligations with respect to the Enterasys software program/firmware installed on the Enterasys product (including any accompanying documentation, hardware or media) ("Program") in the package and prevails over any additional, conflicting or inconsistent terms and conditions appearing on any purchase order or other document submitted by You. "Affiliate" means any person, partnership, corporation, limited liability company, or other form of enterprise that directly or indirectly through one or more intermediaries, controls, or is controlled by, or is under common control with the party specified. This Agreement constitutes the entire understanding between the parties, and supersedes all prior discussions, representations, understandings or agreements, whether oral or in writing, between the parties with respect to the subject matter of this Agreement. The Program may be contained in firmware, chips or other media.

BY INSTALLING OR OTHERWISE USING THE PROGRAM, YOU REPRESENT THAT YOU ARE AUTHORIZED TO ACCEPT THESE TERMS ON BEHALF OF THE END USER (IF THE END USER IS AN ENTITY ON WHOSE BEHALF YOU ARE AUTHORIZED TO ACT, "YOU" AND "YOUR" SHALL BE DEEMED TO REFER TO SUCH ENTITY) AND THAT YOU AGREE THAT YOU ARE BOUND BY THE TERMS OF THIS AGREEMENT, WHICH INCLUDES, AMONG OTHER PROVISIONS, THE LICENSE, THE DISCLAIMER OF WARRANTY AND THE LIMITATION OF LIABILITY. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT OR ARE NOT AUTHORIZED TO ENTER INTO THIS AGREEMENT, ENTERASYS IS UNWILLING TO LICENSE THE PROGRAM TO YOU AND YOU AGREE TO RETURN THE UNOPENED PRODUCT TO ENTERASYS OR YOUR DEALER, IF ANY, WITHIN TEN (10) DAYS FOLLOWING THE DATE OF RECEIPT FOR A FULL REFUND.

IF YOU HAVE ANY QUESTIONS ABOUT THIS AGREEMENT, CONTACT ENTERASYS NETWORKS, LEGAL DEPARTMENT AT (978) 684-1000.

You and Enterasys agree as follows:

1. **LICENSE.** You have the non-exclusive and non-transferable right to use only the one (1) copy of the Program provided in this package subject to the terms and conditions of this Agreement.

2. **RESTRICTIONS.** Except as otherwise authorized in writing by Enterasys, You may not, nor may You permit any third party to:

   (i) Reverse engineer, decompile, disassemble or modify the Program, in whole or in part, including for reasons of error correction or interoperability, except to the extent expressly permitted by applicable law and to the extent the parties shall not be permitted by that applicable law, such rights are expressly excluded. Information necessary to achieve interoperability or correct errors is available from Enterasys upon request and upon payment of Enterasys' applicable fee.

   (ii) Incorporate the Program, in whole or in part, in any other product or create derivative works based on the Program, in whole or in part.

   (iii) Publish, disclose, copy, reproduce or transmit the Program, in whole or in part.

   (iv) Assign, sell, license, sublicense, rent, lease, encumber by way of security interest, pledge or otherwise transfer the Program, in whole or in part.

   (v) Remove any copyright, trademark, proprietary rights, disclaimer or warning notice included on or embedded in any part of the Program.

3. **APPLICABLE LAW.** This Agreement shall be interpreted and governed under the laws and in the state and federal courts of the Commonwealth of Massachusetts without regard to its conflicts of laws provisions. You accept the personal jurisdiction and venue of the Commonwealth of Massachusetts courts. None of the 1980 United Nations Convention on Contracts for the International Sale of Goods, the United Nations Convention on the Limitation Period in the International Sale of Goods, and the Uniform Computer Information Transactions Act shall apply to this Agreement.

4. **EXPORT RESTRICTIONS.** You understand that Enterasys and its Affiliates are subject to regulation by agencies of the U.S. Government, including the U.S. Department of Commerce, which prohibit export or diversion of certain technical products to certain countries, unless a license to export the Program is obtained from the U.S. Government or an exception from obtaining such license may be relied upon by the exporting party.

If the Program is exported from the United States pursuant to the License Exception CIV under the U.S. Export Administration Regulations, You agree that You are a civil end user of the Program and agree that You will use the Program for civil end uses only and not for military purposes.

If the Program is exported from the United States pursuant to the License Exception TSR under the U.S. Export Administration Regulations, in addition to the restriction on transfer set forth in Sections 1 or 2 of this Agreement, You agree not to (i) reexport or release the Program, the source code for the Program or technology to a national of a country in Country Groups D:1 or E:2 (Albania, Armenia, Azerbaijan, Belarus, Bulgaria, Cambodia, Cuba, Estonia, Georgia, Iraq, Kazakhstan, Kyrgyzstan, Laos, Latvia, Libya, Lithuania, Moldova, North Korea, the People's Republic of China, Romania, Russia, Rwanda, Tajikistan, Turkmenistan, Ukraine, Uzbekistan, Vietnam, or such other countries as may be designated by the United States Government), (ii) export to Country Groups D:1 or E:2 (as defined herein) the direct product of the Program or the technology, if such foreign produced direct product is subject to national security controls as identified on the U.S. Commerce Control List, or (iii) if the direct product of the technology is a complete plant or any major component of a plant, export to Country Groups D:1 or E:2 the direct product of the plant or a major component thereof, if such foreign produced direct product is subject to national security controls as identified on the U.S. Commerce Control List or is subject to State Department controls under the U.S. Munitions List.

5. **UNITED STATES GOVERNMENT RESTRICTED RIGHTS.** The enclosed Program (i) was developed solely at private expense; (ii) contains "restricted computer software" submitted with restricted rights in accordance with section 52.227-19 (a) through (d) of the Commercial Computer Software-Restricted Rights Clause and its successors, and (iii) in all respects is proprietary data belonging to Enterasys and/or its suppliers. For Department of Defense units, the Program is considered commercial computer software in accordance with DFARS section 227.7202-3 and its successors, and use, duplication, or disclosure by the Government is subject to restrictions set forth herein.

6. **DISCLAIMER OF WARRANTY.** EXCEPT FOR THOSE WARRANTIES EXPRESSLY PROVIDED TO YOU IN WRITING BY Enterasys, Enterasys DISCLAIMS ALL WARRANTIES, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON- INFRINGEMENT WITH RESPECT TO THE PROGRAM. IF IMPLIED WARRANTIES MAY NOT BE DISCLAIMED BY APPLICABLE LAW, THEN ANY IMPLIED WARRANTIES ARE LIMITED IN DURATION TO THIRTY (30) DAYS AFTER DELIVERY OF THE PROGRAM TO YOU.

7. **LIMITATION OF LIABILITY.** IN NO EVENT SHALL ENTERASYS OR ITS SUPPLIERS BE LIABLE FOR ANY DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS, PROFITS, BUSINESS INTERRUPTION, LOSS OF BUSINESS INFORMATION, SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR RELIANCE DAMAGES, OR OTHER LOSS) ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM, EVEN IF ENTERASYS HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THIS FOREGOING LIMITATION SHALL APPLY REGARDLESS OF THE CAUSE OF ACTION UNDER WHICH DAMAGES ARE SOUGHT.

THE CUMULATIVE LIABILITY OF ENTERASYS TO YOU FOR ALL CLAIMS RELATING TO THE PROGRAM, IN CONTRACT, TORT OR OTHERWISE, SHALL NOT EXCEED THE TOTAL AMOUNT OF FEES PAID TO ENTERASYS BY YOU FOR THE RIGHTS GRANTED HEREIN.

8. **AUDIT RIGHTS.** You hereby acknowledge that the intellectual property rights associated with the Program are of critical value to Enterasys and, accordingly, You hereby agree to maintain complete books, records and accounts showing (i) license fees due and paid, and (ii) the use, copying and deployment of the Program. You also grant to Enterasys and its authorized representatives, upon reasonable notice, the right to audit and examine during Your normal business hours, Your books, records, accounts and hardware devices upon which the Program may be deployed to verify compliance with this Agreement, including the verification of the license fees due and paid Enterasys and the use, copying and deployment of the Program. Enterasys' right of examination shall be exercised reasonably, in good faith and in a manner calculated to not unreasonably interfere with Your business. In the event such audit discovers non-compliance with this Agreement, including copies of the Program made, used or deployed in breach of this Agreement, You shall promptly pay to Enterasys the appropriate license fees. Enterasys reserves the right, to be exercised in its sole discretion and without prior notice, to terminate this license, effective immediately, for failure to comply with this Agreement. Upon any such termination, You shall immediately cease all use of the Program and shall return to Enterasys the Program and all copies of the Program.

9. **OWNERSHIP.** This is a license agreement and not an agreement for sale. You acknowledge and agree that the Program constitutes trade secrets and/or copyrighted material of Enterasys and/or its suppliers. You agree to implement reasonable security measures to protect such trade secrets and copyrighted material. All right, title and interest in and to the Program shall remain with Enterasys and/or its suppliers. All rights not specifically granted to You shall be reserved to Enterasys.

10. **ENFORCEMENT.** You acknowledge and agree that any breach of Sections 2, 4, or 9 of this Agreement by You may cause Enterasys irreparable damage for which recovery of money damages would be inadequate, and that Enterasys may be entitled to seek timely injunctive relief to protect Enterasys' rights under this Agreement in addition to any and all remedies available at law.

11. **ASSIGNMENT.** You may not assign, transfer or sublicense this Agreement or any of Your rights or obligations under this Agreement, except that You may assign this Agreement to any person or entity which acquires substantially all of Your stock or assets. Enterasys may assign this Agreement in its sole discretion. This Agreement shall be binding upon and inure to the benefit of the parties, their legal representatives, permitted transferees, successors and assigns as permitted by this Agreement. Any attempted assignment, transfer or sublicense in violation of the terms of this Agreement shall be void and a breach of this Agreement.

12. **WAIVER.** A waiver by Enterasys of a breach of any of the terms and conditions of this Agreement must be in writing and will not be construed as a waiver of any subsequent breach of such term or condition. Enterasys' failure to enforce a term upon Your breach of such term shall not be construed as a waiver of Your breach or prevent enforcement on any other occasion.

13. **SEVERABILITY.** In the event any provision of this Agreement is found to be invalid, illegal or unenforceable, the validity, legality and enforceability of any of the remaining provisions shall not in any way be affected or impaired thereby, and that provision shall be reformed, construed and enforced to the maximum extent permissible. Any such invalidity, illegality or unenforceability in any jurisdiction shall not invalidate or render illegal or unenforceable such provision in any other jurisdiction.

14. **TERMINATION.** Enterasys may terminate this Agreement immediately upon Your breach of any of the terms and conditions of this Agreement. Upon any such termination, You shall immediately cease all use of the Program and shall return to Enterasys the Program and all copies of the Program.

Contents

Chapter 1: Features Overview

Matrix Series Features 1-1
Factory Default Settings 1-1
Device Management Methods 1-6
Matrix DFE Series and N-SA CLI Overview 1-6

Chapter 2: Getting Started

Starting the Command Line Interface 2-1
Default CLI Access Modes 2-1
Using a Console Port Connection 2-2
Logging in with a Default User Account 2-2
Logging in with an Administratively-Configured User Account 2-2
Using a Telnet Connection 2-3
Setting an IP Address and Basic System Information 2-4
Setting and Verifying a System IP Address 2-4
Setting the Time, Date, and System Information 2-4
Setting System Passwords 2-4
Example 2-5
Setting Syslog Parameters 2-5
About the Matrix DFE-Platinum Series Distributed Console Model 2-6
About Redundant Management on Matrix DFE-Gold Series Modules 2-7
Managing Configuration and Image Files 2-7
Displaying and Changing the Current Configuration 2-8
Downloading a New Firmware Image 2-8
Downloading from an FTP or TFTP Server 2-9
Downloading Using the Serial Port 2-10
Reviewing and Selecting a Boot Firmware Image 2-11
Resetting the Device 2-12
Resetting a Module or the System Immediately 2-12
Scheduling a System Reset 2-12
Displaying Scheduled Reset Information 2-12
Clearing User-Defined Configuration Parameters 2-13
Activating Licensed Features 2-13
Enabling Advanced Capacities on Matrix DFE Series or N-SA Devices 2-13
Using WebView 2-14
Changing WebView Status and Default Port 2-14

Chapter 3: Using the CLI

Configuring CLI Access Security 3-1
Recommendations for Strong, Secure Passwords 3-1
Managing System Passwords and User Accounts 3-3
Starting and Configuring Telnet 3-7
Enabling or Disabling Telnet 3-7
Connecting to a Remote Host Using Telnet 3-7
Configuring Secure Shell (SSH) Server 3-8
Understanding the SSHv2 Protocol 3-8
Setting CLI Parameters 3-9
Displaying Scrolling Screens 3-10

Getting Help with CLI Commands 3-10
Using Context-Sensitive Help 3-10
Performing Keyword Lookups 3-11
Abbreviating and Completing Commands 3-12
Abbreviating a Command 3-12
Using the Spacebar Auto Complete Function 3-12
Using Basic Line Editing Commands 3-13
Understanding Module and Port Numbering 3-13
Console Ports 3-13
Switch Ports 3-14
Designating Ports in the CLI 3-14
Preparing the Device for Router Mode 3-16
Pre-Routing Configuration Tasks 3-16
Enabling Router Configuration Modes 3-17

Chapter 4: Configuring Link Aggregation

Link Aggregation Control Protocol (LACP) Overview 4-1
LACP Operation 4-2
LACP Terminology 4-2
Matrix DFE Series and N-SA Usage Considerations 4-3
Configuring LAG Aggregator Keys and Priority 4-4
Reviewing the Configuration 4-4
Assigning a LAG Aggregator Key 4-5
Assigning a LAG System Priority 4-5
Configuring Underlying Physical Ports 4-5
Statically Assigning Ports to a LAG 4-5
Setting Underlying LACP Port Parameters 4-6
Enabling LACP Flow Regeneration 4-7

Chapter 5: Configuring Spanning Trees

Overview of Spanning Tree Protocols 5-1
Spanning Tree Terms and Definitions 5-1
Spanning Tree (IEEE 802.1D) 5-3
Rapid Spanning Tree (IEEE 802.1w) 5-3
Multiple Spanning Trees (IEEE 802.1s) 5-3
Configuring STP and RSTP 5-8
Reviewing and Re-enabling Spanning Tree 5-8
Adjusting Spanning Tree Parameters 5-9
Enabling the Backup Root Function 5-12
Adjusting RSTP Parameters 5-12
Configuring MSTP 5-14
Simple MSTP Configuration 5-14
Adjusting MSTP Parameters 5-15
Monitoring MSTP 5-15
Configuring Spanguard 5-16
Overview of the Spanguard Function 5-16
Enabling and Adjusting Spanguard 5-17

Chapter 6: Managing Syslog

Logging Overview 6-1
Syslog Terms and Definitions 6-2
Interpreting Messages 6-2
Configuring Syslog Servers, Applications, and Console Logging 6-3
Enabling Syslog Server(s) 6-4
Modifying Syslog Server Defaults 6-4
Reviewing and Configuring Logging for Applications 6-5
Enabling Console Logging and File Storage 6-7

Chapter 7: Configuring IP

Using the Matrix DFE Series or N-SA Device as a Router 7-1
Example Scenario 7-2
Preparing the Device for Router Mode 7-2
Pre-Routing Configuration Tasks 7-3
Enabling Router Configuration Modes 7-3
Reviewing and Configuring Router Interfaces 7-5
Displaying Interface Statistics and Settings 7-5
Configuring Interfaces for IP Routing 7-7
Sample Configuration 7-8
Managing Router Configuration Files 7-10
Displaying the Running Configuration 7-10
Saving or Erasing the Running Configuration 7-11
Removing the Routing Configuration from the Device 7-11
Performing a Basic Router Configuration 7-11
Reviewing and Configuring the ARP Table 7-12
Displaying the ARP Table 7-12
Adding or Removing Static ARP Entries 7-13
Disabling or Re-enabling Proxy ARP on an Interface 7-14
Assigning a MAC Address to an Interface 7-14
Changing the ARP Timeout 7-14
Clearing the ARP Cache 7-14
Configuring IP Broadcast Settings 7-14
Enabling or Disabling IP Directed Broadcasts 7-15
Configuring UDP Broadcast Forwarding 7-15
Configuring Routes and Monitoring IP Traffic 7-16
Adding or Removing Static IP Routes 7-17
Displaying IP Traffic and Performance 7-17
Clearing IP Traffic Counters 7-18
Configuring ICMP 7-18
Enabling or Disabling ICMP 7-18
Sending Ping Requests 7-18
Using Traceroute 7-19

Chapter 8: Configuring OSPF

Using OSPF on a Matrix DFE or N-SA Series Device 8-1
OSPF Overview 8-1
OSPF Terminology 8-2
Supported Functions 8-2
Link-State Advertisements (LSAs) 8-3
Matrix DFE Series and N-SA Operation in OSPF Networks 8-3
OSPF Areas 8-4
Static Multipath Forwarding 8-5
Designating an OSPF Instance 8-5
Virtual Links 8-5
Autonomous System External (ASE) Link Advertisements 8-6
Router Modes Used for OSPF Configuration 8-6
Configuring OSPF 8-7
Modifying Default Configuration Settings 8-7
Activating Advanced Routing 8-9
Performing a Basic OSPF Configuration 8-9
Configuring an OSPF NSSA 8-11
Monitoring and Maintaining OSPF 8-12
Displaying OSPF Information 8-13
Limiting Database Overflow 8-13
Resetting OSPF 8-14
Debugging OSPF 8-14

Chapter 9: Configuring VRRP

VRRP Overview 9-1
VRRP Terms and Definitions 9-1
Configuring VRRP 9-2
Basic VRRP Configuration 9-3
Symmetrical Configuration 9-5
Multi-Backup Configuration 9-8
Modifying a Configuration 9-14
Setting the Backup Priority 9-14
Setting the Advertisement Interval 9-15
Setting a Critical IP Address 9-15
Setting Pre-empt Mode 9-15
Setting an Authentication Key 9-16
Monitoring VRRP 9-17
VRRP Configuration Notes 9-17
Determining the Backup to Master Time Interval 9-17
When the Master Reboots or Is Down 9-18
Determining a Virtual MAC Address 9-18

Figures

2-1 Sample Matrix Startup Screen 2-3
5-1 Example of an MST Region 5-4
5-2 MSTI 1 in a Region 5-6
5-3 MSTI 2 in the Same Region 5-6
5-4 Example of Multiple Regions and MSTIs 5-7
5-5 MSTP Sample Network Configuration 5-15
8-1 OSPF Topology 8-4
8-2 Topology for Basic OSPF Configuration 8-10
9-1 Basic VRRP Configuration 9-3
9-2 Symmetrical VRRP Configuration 9-6
9-3 Multi-Backup VRRP Configuration 9-9

Tables

1-1 Default Device Settings for Basic Switch Operation 1-1
1-2 Default Device Settings for Router Mode Operation 1-4
2-1 CLI Access Modes 2-2
2-2 Advanced License Capacities per Product 2-13
3-1 Password Combinations 3-2
3-2 Default Password Policies 3-6
3-3 CLI Parameters 3-9
3-4 Basic Line Editing Commands 3-13
3-5 Examples of Port String Designations 3-15
3-6 Router Configuration Modes 3-17
4-1 LACP Terms and Definitions 4-2
4-2 Link Aggregation Configurations by Device 4-3
5-1 Spanning Tree Terms and Definitions 5-1
5-2 MSTI Characteristics for Figure 5-4 5-7
5-3 Spanning Tree Port States 5-7
5-4 Commands for Monitoring MSTP 5-15
6-1 Syslog Terms and Definitions 6-2
6-2 Syslog Message Components 6-3
7-1 VLAN and Loopback Interface Configuration Modes 7-2
7-2 Router CLI Configuration Modes 7-3
7-3 show ip interface Command Output 7-6
7-4 show ip arp Output Details 7-14
7-5 Default UDP Forwarding Services 7-16
7-6 Show IP Traffic Commands 7-18
8-1 OSPF Terms and Definitions 8-2
8-2 OSPF Area Restrictions 8-6
8-3 Router Modes Used for OSPF Configuration 8-6
8-4 OSPF Default Settings 8-7
8-5 Displaying OSPF Information 8-13
9-1 VRRP Terms and Definitions 9-1
9-2 Priorities for Virtual Routers Configured on Router 1 9-10
9-3 Priorities for Virtual Routers Configured on Router 2 9-12
9-4 Priorities for Virtual Routers Configured on Router 3 9-13
9-5 VRRP Packet Authentication Field Descriptions 9-16

Procedures

3-1 Enabling the Switch for Routing 3-16
5-1 Configuring Devices 1 and 2 for Simple MSTP 5-15
7-1 Configuring VLANs for Routing 7-8
8-1 Basic OSPF Configuration 8-10
8-2 OSPF NSSA Configuration 8-11
9-1 Configuring Router 1 for Basic VRRP 9-3
9-2 Configuring Router 2 for Basic VRRP 9-4
9-3 Configuring Router 1 for Symmetrical VRRP 9-6
9-4 Configuring Router 2 for Symmetrical VRRP 9-7
9-5 Configuring Router 1 for Multi-Backup VRRP 9-9
9-6 Configuring Router 2 for Multi-Backup VRRP 9-11
9-7 Configuring Router 3 for Multi-Backup VRRP 9-12

1

Features Overview

This chapter provides an overview of the Matrix DFE Series and N-SA devices' unique features and functionality, an overview of the tasks that may be accomplished using the CLI interface, an overview of ways to manage the device, and information on how to contact Enterasys Networks for technical support.

Important Notice

Depending on your Matrix product version (DFE Platinum, DFE Gold, or N Standalone device), and the firmware version used in the device, some features described in this document may not be supported. Refer to the Release Notes shipped with your Matrix device to determine which features are supported.

Matrix Series Features

All Matrix Series chassis-based and standalone modules support secure, business-driven networking with:

- Advanced QoS and policy-based frame classification, and bandwidth management featuring rate limiting, CoS priority queueing, and link aggregation
- Customized, single-source management and control with SNMP, port/VLAN mirroring, Syslog, RMON, multi-image support, and configuration upload/download

In addition, all Matrix DFE modules support:

- New Enterasys advanced Distributed Forwarding Engine (DFE) architecture allowing for single IP address management of an entire chassis
- High availability design featuring fully redundant, hot-swappable modules, power supplies, and fans

Factory Default Settings

The following tables list factory default device settings available on Matrix DFE Series and N-SA devices. Table 1-1 lists default settings for basic switch operation. Table 1-2 lists default settings for router mode operation.

Table 1-1 Default Device Settings for Basic Switch Operation

Device Feature	Default Setting
CDP discovery protocol	Auto enabled on all ports.
CDP authentication code	Set to **00-00-00-00-00-00-00-00**
CDP hold time	Set to **180** seconds.

Table 1-1 Default Device Settings for Basic Switch Operation (continued)

Device Feature	Default Setting
CDP interval	Transmit frequency of CDP messages set to **60** seconds.
Community name	Public.
EAPOL	Disabled.
EAPOL authentication mode	When enabled, set to **auto** for all ports.
GARP timer	Join timer set to **20** centiseconds; leave timer set to **60** centiseconds; leaveall timer set to **1000** centiseconds.
GVRP	Globally enabled.
IGMP	Disabled. When enabled, query interval is set to **125** seconds and response time is set to **100** tenths of a second.
IP mask and gateway	Subnet mask set to **255.0.0.0**; default gateway set to **0.0.0.0**
IP routes	No static routes configured.
Jumbo frame support	Disabled on all ports.
Link aggregation (LACP)	Enabled on all ports.
Link aggregation admin key	Set to **32768** for all ports.
Link aggregation flow regeneration	Disabled.
Link aggregation system priority	Set to **32768** for all ports.
Link aggregation outport algorithm	Set to **DIP-SIP**.
Lockout	Set to disable Read-Write and Read-Only users, and to lockout the default admin (Super User) account for **15** minutes, after **3** failed login attempts,
Logging	Syslog port set to UDP port number **514**. Logging severity level set to **6** (significant conditions) for all applications.
MAC aging time	Set to **300** seconds.
MAC locking	Disabled (globally and on all ports).
MTU discovery protocol	Enabled.
Passwords	Set to an empty string for all default user accounts. User must press ENTER at the password prompt to access CLI.
Password aging	Disabled.
Password history	No passwords are checked for duplication.
Policy classification	Classification rules are automatically enabled when created.
Port auto-negotiation	Enabled on all ports.
Port advertised ability	Maximum ability advertised on all ports.
Port broadcast suppression	Disabled (no broadcast limit).

Table 1-1 Default Device Settings for Basic Switch Operation (continued)

Device Feature	Default Setting
Port duplex mode	Set to **half** duplex, except for 100BASE-FX and 1000BASE-X, which is set to **full** duplex.
Port enable/disable	Enabled.
Port priority	Set to **1**.
Port speed	Set to **10** Mbps, except for 1000BASE-X, which is set to **1000** Mbps, and **100**BASE-FX, which is set to 100 Mbps.
Port trap	All ports are enabled to send link traps.
Priority classification	Classification rules are automatically enabled when created.
RADIUS client	Disabled.
RADIUS last resort action	When the client is enabled, set to **Challenge**.
RADIUS retries	When the client is enabled, set to **3**.
RADIUS timeout	When the client is enabled, set to **20** seconds.
Rate limiting	Disabled (globally and on all ports).
SNMP	Enabled.
SNTP	Disabled.
Spanning Tree	Enabled (globally and on all ports).
Spanning Tree edge port administrative status	Enabled.
Spanning Tree edge port delay	Enabled.
Spanning Tree forward delay	Set to **15** seconds.
Spanning Tree hello interval	Set to **2** seconds.
Spanning Tree ID (SID)	Set to **0**.
Spanning Tree legacy path cost	Disabled.
Spanning Tree maximum aging time	Set to **20** seconds.
Spanning Tree point-to-point	Set to **auto** for all Spanning Tree ports.
Spanning Tree port priority	All ports with bridge priority are set to **128** (medium priority).
Spanning Tree priority	Bridge priority is set to **32768**.
Spanning Tree topology change trap suppression	Enabled.
Spanning Tree transmit hold count	Set to **3**.

Table 1-1 Default Device Settings for Basic Switch Operation (continued)

Device Feature	Default Setting
Spanning Tree version	Set to **mstp** (Multiple Spanning Tree Protocol).
SSH	Disabled.
System baud rate	Set to **9600** baud.
System contact	Set to empty string.
System location	Set to empty string.
System name	Set to empty string.
Terminal	CLI display set to **80** columns and **24** rows.
Timeout	Set to **15** minutes.
User names	Login accounts set to **ro** for Read-Only access; **rw** for Read-Write access; and **admin** for Super User access.
VLAN dynamic egress	Disabled on all VLANs.
VLAN ID	All ports use a VLAN identifier of **1**.

Table 1-2 Default Device Settings for Router Mode Operation

Device Feature	Default Setting
Access groups (IP security)	None configured.
Access lists (IP security)	None configured.
Area authentication (OSPF)	Disabled.
Area default cost (OSPF)	Set to **1**.
Area NSSA (OSPF)	None configured.
Area range (OSPF)	None configured.
ARP table	No permanent entries configured.
ARP timeout	Set to **14,400** seconds.
Authentication key (RIP and OSPF)	None configured.
Authentication mode (RIP and OSPF)	None configured.
Dead interval (OSPF)	Set to **40** seconds.
Disable triggered updates (RIP)	Triggered updates allowed.
Distribute list (RIP)	No filters applied.
DoS prevention	Disabled.
DVMRP	Disabled. Metric set to **1**.
Hello interval (OSPF)	Set to **10** seconds for broadcast and point-to-point networks. Set to **30** seconds for non-broadcast and point-to-multipoint networks.

Table 1-2 Default Device Settings for Router Mode Operation (continued)

Device Feature	Default Setting
ICMP	Enabled for echo-reply and mask-reply modes.
IP-directed broadcasts	Disabled.
IP forward-protocol	Enabled with no port specified.
IP interfaces	Disabled with no IP addresses specified.
IRDP	Disabled on all interfaces. When enabled, maximum advertisement interval is set to **600** seconds, minimum advertisement interval is set to **450** seconds, holdtime is set to **1800** seconds, and address preference is set to **0**.
MD5 authentication (OSPF)	Disabled with no password set.
MTU size	Set to **1500** bytes on all interfaces.
OSPF	Disabled.
OSPF cost	Set to **10** for all interfaces.
OSPF network	None configured.
OSPF priority	Set to **1**.
Passive interfaces (RIP)	None configured.
Proxy ARP	Enabled on all interfaces.
Receive interfaces (RIP)	Enabled on all interfaces.
Retransmit delay (OSPF)	Set to **1** second.
Retransmit interval (OSPF)	Set to **5** seconds.
RIP receive version	Set to accept both version **1** and version **2**.
RIP send version	Set to version **1**.
RIP offset	No value applied.
SNMP	Enabled.
Split horizon	Enabled for RIP packets without poison reverse.
Stub area (OSPF)	None configured.
Telnet	Enabled.
Telnet port (IP)	Set to port number **23**.
Timers (OSPF)	SPF delay set to **5** seconds. SPF holdtime set to **10** seconds.
Transmit delay (OSPF)	Set to **1** second.
VRRP	Disabled.

Device Management Methods

The Matrix DFE Series and N-SA devices can be managed using the following methods:

- Locally using a VT type terminal connected to the console port.
- Remotely using a VT type terminal connected through a modem.
- Remotely using an SNMP management station.
- In-band through a Telnet connection.
- In-band using Enterasys Networks' NetSight® management application.
- Remotely using WebView™, Enterasys Networks' embedded web server application.

For detailed setup instructions for connecting a terminal or modem to the Matrix device, refer to the Matrix DFE Series or N-SA Installation Guide for your product.

Matrix DFE Series and N-SA CLI Overview

Enterasys Networks' Matrix DFE Series and N-SA CLI interface allows you to perform a variety of network management tasks, including the following:

- Assign IP address and subnet mask.
- Select a default gateway.
- Assign login passwords and user accounts for additional security.
- Download a new firmware image.
- Designate which network management workstations receive SNMP traps from the device.
- View device, interface, and RMON statistics.
- Manage configuration files.
- Assign ports to operate in the standard or full duplex mode.
- Control the number of received broadcasts that are switched to the other interfaces.
- Set flow control on a port-by-port basis.
- Set port configurations and port-based VLANs.
- Configure ports to prioritize and assign a VLAN or Class of Service to incoming frames based on Layer 2, Layer 3, and Layer 4 information.
- Configure the device to operate as a Generic Attribute Registration Protocol (GARP) device to dynamically create VLANs across a switched network.
- Redirect frames according to a port or VLAN and transmit them on a preselected destination port.
- Configure Multiple Spanning Trees.
- Clear NVRAM.
- Configure interfaces for IP routing.
- Configure RIP, OSPF, DVMRP, IRDP, and VRRP routing protocols.
- Configure security, including 802.1X. RADIUS, SSHv2, MAC locking, MAC authentication, and DoS attack prevention.
- Configure access lists (ACLs).

2

Getting Started

This chapter provides information about the following basic setup procedures on the Matrix DFE Series / N-SA device.

For information about...	Refer to page...
Starting the Command Line Interface	2-1
Setting an IP Address and Basic System Information	2-4
Setting System Passwords	2-4
Setting Syslog Parameters	2-5
About the Matrix DFE-Platinum Series Distributed Console Model	2-6
About Redundant Management on Matrix DFE-Gold Series Modules	2-7
Managing Configuration and Image Files	2-7
Downloading a New Firmware Image	2-8
Reviewing and Selecting a Boot Firmware Image	2-11
Resetting the Device	2-12
Activating Licensed Features	2-13
Using WebView	2-14

Starting the Command Line Interface

Default CLI Access Modes

The Matrix DFE Series / N-SA Command Line Interface (CLI) provides three default access modes. Table 2-1 shows these default access modes, the corresponding user account names, the command prompt the CLI displays, and the privileges associated with each mode. For more information on changing these settings, refer to "Setting System Passwords" on page 2-4, and "Managing System Passwords and User Accounts" on page 3-3.

Note: Depending on which Matrix Series device you are using, your default command prompt may be different than the examples shown.

Table 2-1 CLI Access Modes

Mode (User Name)	Prompt	Access Privileges
Read-Only (**ro**)	`Matrix(ro)->`	Permitted to view Read-Only (**show**) commands.
Read-Write (**rw**)	`Matrix(rw)->`	Permitted to modify all modifiable parameters in **set** and **show** commands, as well as view Read-Only commands.
Admin / Super User (**admin**)	`Matrix(su)->`	Permitted all Read-Write and Read-Only privileges, and the ability to modify local user accounts.

Note: Unless otherwise specified, the instructions in this section refer to entering commands in Read-Write (**rw**) access mode.

Using a Console Port Connection

Once you have connected a terminal to the local console port as described in the Matrix DFE Series / N-SA Installation Guide for your product, the startup screen, as shown in the example in Figure 2-1, will display. You can now start the CLI as described in this section by:

- Logging in with a Default User Account or
- Logging in with an Administratively-Configured User Account

Note: By default, the Matrix system password is set to a blank string. For information on changing these default settings, refer to Setting System Passwords on page 2-4.

Logging in with a Default User Account

If this is the first time you are logging in to the Matrix device, or if the default user accounts have not been administratively changed, proceed as follows:

1. At the login prompt, enter one of the following default user names:
   - **ro** for Read-Only access,
   - **rw** for Read-Write access.
   - **admin** for Super User access.
2. Press ENTER. The Password prompt displays.
3. Leave this string blank and press ENTER. The device information and Matrix prompt displays as shown in the example in Figure 2-1.

Logging in with an Administratively-Configured User Account

If the device's default user account settings have been changed, proceed as follows:

1. At the login prompt, enter your administratively-assigned user name and press ENTER.
2. At the Password prompt, enter your password and press ENTER.

The device information and Matrix prompt displays as shown in the example in Figure 2-1.

Note: Users with Read-Write (**rw**) and Read-Only (**ro**) access can use the **set password** command to change their own passwords. Administrators with Super User (**su**) access can use the **set system login** command to create and change user accounts, and the **set password** command to change any local account password. For more information about creating and managing user accounts and passwords, refer to "Managing System Passwords and User Accounts" on page 3-3.

Using a Telnet Connection

Once the Matrix DFE Series / N-SA device has a valid IP address, you can establish a Telnet session from any TCP/IP based node on the network as follows.

1. Telnet to the device's IP address.
2. Enter login (user name) and password information in one of the following ways:
   - If the device's default login and password settings have not been changed, follow the steps listed in "Logging in with a Default User Account" on page 2-2, or
   - Enter an administratively-configured user name and password.

The device information and the Matrix prompt displays as shown in the example in Figure 2-1.

For information about setting the IP address, refer to "Setting and Verifying a System IP Address" on page 2-4.

For information about configuring Telnet settings, refer to "Starting and Configuring Telnet" on page 3-7.

Refer to the instructions included with the Telnet application for information about establishing a Telnet session.

Figure 2-1 Sample Matrix Startup Screen

```
login: rw
Password:


MATRIX N7 PLATINUM
Command Line Interface

Enterasys Networks, Inc.
50 Minuteman Rd.
Andover, MA 01810-1008 U.S.A.

Phone:  +1 978 684 1000
E-mail: support@enterasys.com
WWW:    http://www.enterasys.com

(c) Copyright Enterasys Networks, Inc. 2005

Module Serial Number:     TRI-A26
Module Firmware Revision: 05.11.00

Matrix N7 Platinum(rw)->
```

Setting an IP Address and Basic System Information

Use the procedures in this section to set an IP address for the system and to set basic system information, including the date and time the system will display, where the system is located and a system contact person within your organization.

Note: Some of the commands in these procedures accept a string value. String values can be up to a maximum of 255 characters in length, including blank spaces. Surround strings that contain blanks with quotation marks (Example: **"string with internal blanks"**).

Setting and Verifying a System IP Address

Use the following commands to set and verify the IP address:

1. Set the system IP address, optional subnet mask, and optional default gateway:

   set ip address *ip-address* [**mask** *ip-mask*] [**gateway** *ip-gateway*]

   If not specified, *ip-mask* will be set to the natural mask of the *ip-address* and *ip-gateway* will be set to the *ip-address*.

2. (Optional) Verify the system IP address:

   show ip address

Setting the Time, Date, and System Information

Use the following commands to set the system time, name, location, and contact:

1. Set the system time in a month, day, year, and/or a 24-hour format:

   set time [*mm/dd/yyyy*] [*hh:mm:ss*]

2. Set a system name:

   set system name [*string*]

3. Set a system location:

   set system location [*string*]

4. Set a contact name (the person to contact regarding this device):

   set system contact [*string*]

Setting System Passwords

By default, the Matrix DFE Series or N-SA device provides three access modes, or account "names". These are Admin or Super-User (**su**), Read-Write (**rw**), and Read-Only (**ro**). At startup, no passwords are defined. Administrators with Super-User access can set a new system password with the following commands.

1. To set a default system password, from the **su** prompt enter:

   set password {admin|**rw** |**ro**}

2. Press ENTER. When prompted, enter a password as shown in the example below. Passwords are case sensitive and must be a minimum of 8 characters and a maximum of 40 characters.

Example

This example shows how an administrator would change the Read-Write password from the system default (blank string):

```
Matrix(su)->set password rw
Please enter new password: ********
Please re-enter new password: ********
Password changed.
Matrix(su)->
```

Caution: Test new passwords before saving the active configuration to the startup configuration file. To keep your passwords secure, the Matrix DFE Series / N-SA device does not have a command for displaying passwords.

If you forget a Read-Write or Read-Only password, you can remove it by entering the following command at the **su** prompt:

clear system login {**rw** | **ro**}

If you forget the Admin password, it can be reset by toggling dip switch 8 on the device as described in the the Matrix DFE Series / N-SA Installation Guide for your product.

For more information, refer to "Configuring CLI Access Security" on page 3-1.

Setting Syslog Parameters

The CLI can use Syslog messages to communicate eight levels of system error messages to a Syslog server. Valid levels and their corresponding severity designations are:

1 - emergencies (system is unusable)

2 - alerts (immediate action required)

3 - critical conditions

4 - error conditions

5 - warning conditions (default level)

6 - notifications (significant conditions)

7 - informational messages

8 - debugging messages

The Matrix DFE Series or N-SA device writes the Syslog messages to a Syslog daemon on UDP port 514. By default, the CLI sends level 6 messages about all system applications to the specified Syslog server. You can use the commands in this section to specify a Syslog server and set the CLI to send all or only some of the message types.

1. Add a Syslog server to the Matrix DFE Series or N-SA device, and change the message severity level:

   set logging server *index* **ip-addr** *ip-addr* **severity** *severity* **enable**

Example

This example shows how to enable Syslog server 1 at IP address 134.141.89.113 to log messages at severity level 3 (critical conditions only):

```
Matrix(rw)->set logging server 1 ip-addr 134.141.89.113 severity 3 enable
```

2. (Optional) Verify the Syslog settings.

   show logging server *index*

For more information, refer to Chapter 6, "Managing Syslog".

About the Matrix DFE-Platinum Series Distributed Console Model

Note: The distributed console and redundant management features described in this section do not apply to Matrix DFE-Gold or Matrix N-SA (standalone) devices.

Because each Matrix DFE-Platinum module is part of a system, one module is elected to control system management functions. This System Management Module (SMM) will coordinate and control the configuration of the entire chassis. Access to the SMM is available through any console (COM) port on any module in the chassis. Only one local console CLI session can be active at any one time, and active status is granted to the first connection to any of the console ports. SMM status is indicated when the MGMT LED on the module's faceplate is solid green. The following section provides an overview of the SMM console selection process.

The System Management Module (SMM) Console

The SMM console is simply a CLI session that originates from the SMM module in any given chassis. This CLI session is the only one that will accept CLI commands and is, therefore, the only serial interface into the CLI command set. Although access to the SMM console is available from the console port of any module in the Matrix DFE-Platinum chassis, control is granted to only one physical connection at a time. The active connection is determined on a first-come-first-serve basis with the first rights always given to the physical connection of the SMM console, if present.

Accessing the CLI Through a Console Port

With one serial connection to a console port as described in the *Matrix DFE-Platinum Series Installation Guide*, CLI access will always be granted to the SMM console.

With more than one physical connection, one active session to the SMM console will be allowed, and the other physical connections will display a message indicating that the SMM console is in use. On device startup, the CLI login screen is generally displayed on all of the console ports at first. Once an SMM module has been elected, only one CLI session will stay active and messages will be printed to the other modules indicating that the SMM console is in use elsewhere.

The CLI session will run on the SMM module regardless of which console port has been connected. If **exit** is entered or a CLI timeout occurs, the CLI and login process will simply restart. This will not relinquish control of the SMM console. Control of the SMM console may be lost, however, if the serial connection is removed or if the module is reset.

For more information on connecting a terminal or modem to a console port, refer to the *Matrix DFE-Platinum Series Installation Guide*.

About Redundant Management on Matrix DFE-Gold Series Modules

Notes: Interoperability of Matrix DFE-Gold Series modules is dependent upon module placement rules during installation in the chassis. For details on these rules and their effects on system management, refer to the *Matrix DFE-Gold Series Installation Guide*.

The distributed console and redundant management features described in this section do not apply to Matrix DFE-Platinum or Matrix N-SA (standalone) devices.

The DFE-Gold System Management Module (SMM) coordinates and controls the configuration of the entire chassis. By default, this is the module installed in slot 1. Access to the SMM is available through any console (COM) port on any module in the chassis. Only one CLI session can be active at any one time, and active status is granted to the first connection to any of the console ports.

In order to enable switch and routing redundancy on a Matrix DFE-Gold Series device, you must purchase and activate a license key. If you have purchased a redundancy license, you can proceed to activate it as described in this section. If you wish to purchase a redundancy license, contact Enterasys Networks Sales.

When a redundancy license key is purchased and activated as described in "Activating Licensed Features" on page 2-13, redundancy can be configured on the module in slot 2 of the chassis. Then, in the event module 1 fails, module 2 will assume chassis management.

Managing Configuration and Image Files

Matrix DFE Series and N-SA devices provide a single configuration interface which allows you to perform both switch and router configuration with the same command set. Matrix DFE devices also provide redundant, distributed copies of each image file in the event that DFE modules are added or removed from the Matrix chassis. When managing configuration and image files on Matrix DFE Series devices, it is important to remember the following considerations:

- All modules will have the same image files loaded. If a module is inserted into a chassis, any images unique to that module will be deleted, and any images not on that module will be copied to it.
- Configuration files are stored per slot, with each slot in the chassis carrying unique files with it and not inheriting files.

The following section describes commands for managing both switch and router configuration.

For details on performing a basic routing configuration (while operating in router mode), refer to "Performing a Basic Router Configuration" on page 7-11.

For details on downloading a new firmware image, refer to page 2-8.

For details on reviewing and selecting the boot firmware image, refer to page 2-11.

Note: The commands described in this section manage both switch and router configuration parameters, but must be executed from the switch CLI.

Displaying and Changing the Current Configuration

Use these commands to display and change the current configuration:

1. Display current default and non-default configuration settings:

   show config all

2. (Optional) View one or all configuration files stored in the system:

   dir [*filename*]

3. (Optional) View the contents of a specific configuration file:

   show file *filename*

4. Upload or download a new configuration file.

   copy *source destination*

Notes: The Matrix DFE module to which a configuration file is downloaded must have the same hardware configuration as the Matrix DFE module from which it was uploaded.

A **slot1** designation must be entered, as shown in the examples below, when specifying configuration files on the Matrix N-SA device.

Examples

This example shows how to download a configuration file using TFTP to the Matrix N-SA device or the DFE module in slot 1:

```
Matrix(rw)->copy tftp://134.141.89.34/myconfig slot1/myconfig
```

This example shows how to upload a configuration file using Anonymous FTP from the Matrix N-SA device or the DFE module in slot 1:

```
Matrix(rw)->copy slot1/myconfig ftp://134.141.89.34/myconfig
```

This example shows how to copy a configuration file from the DFE module in slot 3 to the DFE module in slot 5:

```
Matrix(rw)->copy slot3/myconfig slot5/myconfig
```

5. Execute a previously downloaded configuration file stored on the device.

   configure *filename* [**append**]

   Specifying **append** will execute the configuration as an appendage to the current running configuration. This is equivalent to typing the contents of the config file directly into the CLI and can be used, for example, to make incremental adjustments to the current configuration. If **append** is not specified, the current running configuration will be replaced with the contents of the specified configuration file, which will require an automated reset of the chassis

For more information, refer to "Saving or Erasing the Running Configuration" (page 7-11).

Downloading a New Firmware Image

You can upgrade the operational firmware in the Matrix DFE Series and N-SA devices without physically opening the device or being in the same location. There are two ways to download firmware to the device:

- Using FTP download. This procedure uses an FTP server connected to the network and downloads the firmware using the FTP protocol. It is the most robust downloading mechanism.
- Using TFTP download. This procedure uses a TFTP server connected to the network and downloads the firmware using the TFTP protocol.
- Using the serial (console) port. This procedure is an out-of-band operation that copies the firmware through the serial port to the device. It takes approximately five minutes and requires minimal configuration. It should be used in cases when you cannot connect the device to perform the in-band download procedure using FTP or TFTP. Serial console download has been successfully tested with the following applications:
  - HyperTerminal Copyright 1999
  - Tera Term Pro Version 2.3

Any other terminal applications may work but are not explicitly supported. For details, refer to "Downloading Using the Serial Port" on page 2-10.

Important Notice

Matrix DFE Series and N-SA devices allow you to download and store multiple image files. This feature is useful for reverting back to a previous version in the event that a firmware upgrade fails to boot successfully. After downloading firmware as described in this section, you can select which image file you want the device to load at startup using the **setboot** command in the System Image Loader menu, or the **set boot system** command as described on page 2-11.

Use the **show version** command to display your current firmware (**fw**) and BootProm (**bp**) settings.

Downloading from an FTP or TFTP Server

To perform an FTP or TFTP download:

1. If you have not already done so, set the device's IP address as detailed in "Setting and Verifying a System IP Address" (page 2-4).
2. (Optional) View one or all image files stored in the system:

   dir [*filename*]
3. Download a new image file:

   copy *source destination*

Examples

This example shows how to download an image using TFTP:

```
Matrix(rw)->copy tftp://134.141.89.34/ets-mtxe7-msi images/newimage
```

This example shows how to download an image using Anonymous FTP:

```
Matrix(rw)->copy ftp://134.141.89.34/ets-mtxe7-msi images/newimage
```

This example shows how to download an image using FTP with user credentials:

```
Matrix(rw)->copy tftp://user:passwd@134.141.89.34/ets-mtxe7-msi images/
newimage
```

4. You can now set the device to load the new image file at startup using the **set boot system** command as described in "Reviewing and Selecting a Boot Firmware Image" (page 2-11).

Downloading Using the Serial Port

To download device firmware using the serial (console) port:

1. With the console port connected as described in the Matrix DFE Series or N-SA Installation Guide for your product, power up the device. The following message displays:

```
Boot ROM Initialization, Version 01.00.01

Copyright (c) 2005 Enterasys Networks, Inc.

SDRAM size: 128 MB
Testing SDRAM....                       PASSED.
Loading Boot Image: 01.00.02...         DONE.

Uncompressing Boot Image...             DONE.

Press any key to enter System Image Loader menu
```

2. Before the boot up completes, press any key. The System Image Loader prompt displays.
3. (Optional) Type **?** and press ENTER. The following System Image Loader menu displays:

```
[System Image Loader]: ?

 ?, help                        - print this list
 boot                           - boot (load and go)
 download                       - start ZMODEM download
 list                           - display available images
 log                            - message log
 setbaud <rate>                 - set baud rate, (9600,38400,57600,115200)
 setboot <filename>             - change boot image file
 showboot                       - display boot image file
```

4. Type **setbaud 115200** to set the device baud rate to 115200 and press ENTER.
5. Open your terminal application and, from that application, set the terminal baud rate to **115200**.
6. From the System Image Loader prompt, type **download** to start the ZMODEM receive process.
7. Send the image file using the ZMODEM protocol from your terminal application. (This procedure will vary depending on your application.) When the ZMODEM download is finished, the following message displays:

```
[System Image Loader]: download
Preparing to receive file...

Writing file...
Download successful.
[System Image Loader]:
```

8. From the System Image Loader prompt, type **setbaud 9600** to set the device baud rate back to **9600** and press ENTER**.**
9. From your terminal application, set the terminal baud rate back to **9600**.
10. Type **setboot** *filename* to set the device to boot to the new firmware image and press ENTER. In this example, the downloaded image file is named "myimage." The following message displays:

```
[System Image Loader]: setboot myimage
Image boot file set to myimage
[System Image Loader]:
```

11. Type **boot** and press ENTER to reboot the device. The following message indicates the downloaded image booted successfully:

```
[System Image Loader]: boot
/flash0/  - Volume is OK
Loading myimage...                    DONE.
```

The system must be reset by software for a new boot image to take effect at startup. If the chassis is powered OFF and then back ON, the current active image will just reload at startup.

Note: If you reboot without specifying the image to boot with **setboot** as described above, the device will attempt to load whatever image is currently stored in the bootstring using the **set boot system** command. If the device cannot find the image, or it is not set, it will search through available images and attempt to boot the newest one. It will then set the bootstring to whatever image file name was successfully loaded.

Reviewing and Selecting a Boot Firmware Image

Use the following commands to display and change your current boot image.

1. Display the firmware image the switch loads at startup:

   show boot system

2. (Optional) View one or all image files stored in the system:

   dir [*filename*]

   In the command output, "Active" indicates the image that is currently running, and "Boot" indicates the image that is currently scheduled to boot next.

3. Specify a new boot firmware image. This will move the boot designation from the current running image, but will allow the active image to stay where it is until after a device reset , when that image has actually been booted.

   set boot system *filename*

When you reboot the device, you can verify the new boot image with either the **show version** or **show boot system** command.

Resetting the Device

Use the procedures in this section to reset one or more device modules, to schedule a system reset in order to load a new boot image, or to clear the user-defined switch and router configuration.

Resetting a Module or the System Immediately

Use the following command to reset a module or the entire device without losing any user-defined configuration settings:

reset [*mod-num* | **system**]

If *mod-num* is not specified, the entire system will be reset.

Note: A Matrix DFE Series or N-SA module can also be reset with the RESET switch located on its front panel. For information on how to do this, refer to the Matrix DFE Series or N-SA Installation Guide shipped with your module.

Scheduling a System Reset

When loading a new boot image, it is useful to schedule a system reset as follows.

1. To schedule a reset at a specific future time, using a 24-hour format:

   reset at *hh:mm* [*mm/dd*] [*reason*]

   If month and day are not specified, the reset will be scheduled for the first occurrence of the specified time.

   A *reason* text string containing spaces must be enclosed in quotations. If one is not specified for the reset, none will be applied.

2. To schedule a reset after a specified time lapse:

   reset in *hh:mm* [*reason*]

Displaying Scheduled Reset Information

Use the following command to display information about scheduled device resets:

show reset

Example

```
Matrix(rw)->show reset

Reset scheduled for Fri Jan 21 2004, 23:00:00 (in 3 days 12 hours 56 minutes 57
seconds).
Reset reason: Software upgrade
```

Clearing User-Defined Configuration Parameters

Use the following command to clear the user-defined system configuration parameters for one or more modules and reset those modules back to factory defaults:

```
clear config mod-num | all
```

Note: This command does not affect the switch IP address.

If the module being reset is in a chassis with other active modules, it will inherit system settings from the system. For a list of factory device default settings, refer to Table 1-1 and Table 1-2 on page 1-4.

Activating Licensed Features

In order to enable advanced features, such as increased user and port capacity, redundant management, advanced routing protocols, and extended ACLs, you must purchase and activate the license key(s) applicable to your requirements and product as shown in Table 2-2. If you have purchased one or more licenses, you can proceed to activate them as described in this section. If you wish to purchase a license, contact Enterasys Networks Sales.

Table 2-2 Advanced License Capacities per Product

License	Advanced Features	DFE-Platinum	DFE-Gold	N-SA
N-EOS-L3	Advanced routing functions, including OSPF, DVMRP, PIM, LSNAT, and extended ACLs	●	●	●
N-EOS-PPC	Increased per-port user capacity to 256, up to a maximum of 1024 users per module	●		●
N-EOS-PUC	Increased user capacity to 2048 users per chassis	●		
N-EOS-RED	Management module redundancy		●	

Note: Increased per-port user capacity licenses (N-EOS-PPC) must be purchased per slot/module.

Enabling Advanced Capacities on Matrix DFE Series or N-SA Devices

Use the following procedure to activate advanced license capabilities on Matrix DFE Series or N-SA devices:

1. Determine your appropriate *license-key* from the activation key document shipped with your product, or by viewing the top of the **running-config** output:

   ```
   show running-config
   ```

2. Enter the applicable license key for the capability you want to activate:

   ```
   set license {advanced | redundancy | user-capacity | port-capacity} license-key [slot slot]
   ```

3. Activate the license by resetting the device:

   ```
   reset [mod-num | system ]
   ```

4. Verify license activation:

   show license

Using WebView

WebView is the Enterasys Networks embedded web server for device configuration and management tasks. By default, WebView is enabled on TCP port number 80 of the Matrix device.

You can verify WebView status, enable or disable WebView, and reset the WebView port as described in the following section.

Changing WebView Status and Default Port

Use the following commands to display WebView status, enable or disable WebView, or to set the WebView port:

1. (Optional) Verify WebView status:

   show webview

2. Enable or disable WebView:

   set webview {**enable** | **disable**}

3. Set a different TCP port through which to run WebView:

   set webview port *webview_port*

3

Using the CLI

This chapter provides information about using the CLI on Matrix DFE Series / N-SA devices.

For information about...	Refer to page...
Configuring CLI Access Security	3-1
Starting and Configuring Telnet	3-7
Configuring Secure Shell (SSH) Server	3-8
Setting CLI Parameters	3-9
Getting Help with CLI Commands	3-10
Abbreviating and Completing Commands	3-12
Using Basic Line Editing Commands	3-13
Understanding Module and Port Numbering	3-13
Preparing the Device for Router Mode	3-16

For details on starting the CLI, refer to Chapter 2, "Starting the Command Line Interface".

Configuring CLI Access Security

When configuring your network access security policy, Enterasys recommends that you employ at least the following:

- Minimum password length of 8 characters.
- Number of failed login attempts before disabling a user's account should not exceed 6.
- New login attempts cannot be made for at least 60 minutes after disabling a user account.
- Changing administrator-assigned user password after first login (not enabled by default).
- Password lifetime of no more than 90 days.
- Not reusing at least the previous 5 passwords.

Recommendations for Strong, Secure Passwords

Enterasys recommends avoiding the following types of passwords:

- User's name (first or last), child's name, or the name of a pet
- Birthday or anniversary
- "Password"

- Repeated characters (e.g., “AAAAAA” or “999999")
- Sports teams or terms (such as “Bulls” or “Golfer”)
- Favorite recording artist
- Obscenities or sexual terms

In addition to avoiding weak passwords, Enterasys also recommends that you:

- Do NOT write down the password and post it near the terminal.
- Do NOT use the login name and password of a former employee.
- Make sure that someone besides the network administrator knows the administrative (master) account user name and password and tests them periodically. This prevents you from losing access to your network should anything happen to the employee or if the relationship with that employee deteriorates.
- Avoid using the master account for anything but administration. Using this account frequently to perform mundane network operations can lead to unnecessary accidents.
- Disable network “guest” accounts when not in use.

Creating Strong Passwords

Computer programs used to hack passwords fall into the following two general categories:

- **Dictionary attacks** involve using words from a “dictionary” or database of frequently-used terms to attempt to match the user password (for example: names, cities, and sports teams).
- **Brute force attacks** focus on discovering passwords using software that generates a sequence of character combinations. For example, guessing a three-character password that uses only alphabetical characters (A - Z) would involve testing all letter combinations between AAA and ZZZ—a total of 17,576 possible combinations.

A strong password is vital because password cracking tools continue to improve and the computers used to crack passwords are more powerful than ever. Network passwords that once took weeks to break are now be broken in a few hours. The strength of a password is proportional to its length and complexity because the longer the password, the longer it takes to generate all possible combinations. Table 3-1 depicts the relative strength associated with password length and the character types allowed.

Table 3-1 Password Combinations

Number of Characters in Password	Possible Combinations (Letters A-Z only)	Possible Combinations (Letters A-Z, with numbers 0-9)
1	26	36
2	676	1,296
3	17,576	46,656
4	456,976	1,679,616
5	11,881,376	60,466,176
6	308,915,776	2,176,782,336
7	8,031,810,176	78,364,164,096
8	208,827,064,576	2,821,109,907,456

Table 3-1 Password Combinations (continued)

Number of Characters in Password	Possible Combinations (Letters A-Z only)	Possible Combinations (Letters A-Z, with numbers 0-9)
9	5,429,503,678,976	101,559,956,668,416
10	141,167,095,653,376	3,656,158,440,062,980

For a password to be strong and hard to break, it should:

- Be at least 8 characters long.
- Contain characters from each of the following groups:
  - Letters (uppercase and lowercase) — A, B, C; a, b, c
  - Numerals — 0,1,2,3,4,5,6,7,8,9
  - Symbols — ` ~ ! @ # $ % ^ & * ( ) _ + - = { } | [ ] \ : " ; ' < > ? , . /
- Have at least one symbol character in the second through sixth positions.
- Be significantly different from prior passwords.
- Not contain your name or user name.
- Not be a common word or name.

Managing System Passwords and User Accounts

By default, the Matrix DFE Series or N-SA device operates in single-user mode with password access enabled and no passwords defined. It provides three access modes, or account "names." These are Admin or Super-User (**su**), Read-Write (**rw**), and Read-Only (**ro**). For more information about these default access accounts and their associated privileges, refer back to Table 2-1 on page 2-2. This section provides information about the following system password and user account management tasks that can be completed by an administrator with Super User access.

- Setting New System Passwords
- Creating and Managing User Accounts
- Setting Password Policies, including age out time, the number of previously-used passwords checked for duplication when users create new passwords, and the number of failed login attempts allowed before user lock out.

Note: Unless otherwise specified, commands described in this section are restricted to administrators with Super-User (**su**) access.

Once configured, passwords can be changed with the **set password** command. Administrators with Super User access can change any password on the system as described in the following section. Individual users with Read-Write access can change their own passwords, as described on page 3-5, but cannot enter or modify other system passwords. Only users with admin (**su**) access privileges can change any password on the system.

Setting New System Passwords

1. To set system passwords for the default Admin, Read-Write or Read-Only user accounts, from the **su** prompt enter:

   set password *{***admin** | **rw** | **ro***}*

2. Press ENTER. When prompted, enter a password as shown in the example below. Passwords are case sensitive and must be a minimum of 8 characters and a maximum of 40 characters.

Example

This example shows how an administrator would change the Read-Write password from the system default (blank string):

```
Matrix(su)->set password rw
Please enter new password: ********
Please re-enter new password: ********
Password changed.
Matrix(su)->
```

The Matrix DFE Series or N-SA device stores passwords in the startup configuration file. If you copy a configuration file from one Matrix DFE Series or N-SA device to another, the passwords in the file will automatically apply to the new device. When you activate a new password by copying the **set password** command to the active configuration, the Matrix DFE Series or N-SA device hides the password text in the configuration file to prevent others who access this file from viewing the password.

Caution: Test new passwords before saving the active configuration to the startup configuration file. To keep your passwords secure, the Matrix DFE Series or N-SA device does not have a command for displaying passwords.

If you forget a Read-Write or Read-Only password, you can remove it by entering the following command at the **su** prompt:

clear system login {**rw** | **ro***}*

If you forget the Admin password, it can be reset by toggling dip switch 8 on the device as described in the Matrix DFE Series or N-SA Installation Guide for your product.

Creating and Managing User Accounts

Multi-user mode password security employs individual user accounts to grant CLI permissions on a case-by-case basis—and requires that each user log in using user name and password. The Matrix DFE Series or N-SA device supports up to 16 concurrent user accounts, including the admin account, which cannot be disabled or deleted.

Administrators with Super-User access can use the following commands to create, disable, remove and/or verify individual user accounts.

Creating a New Login User Account

From the **su** prompt, use the following command to create a new user account with access privileges:

set system login *username* {**super-user** | **read-write** | **read-only**} **enable**

Applying Access Privileges to an Existing User Account

From the **su** prompt, use the following command to apply access privileges to an existing account:

set system login *username* {**super-user** | **read-write** | **read-only**}

Disabling a User Account

From the **su** prompt, use the following command to disable an existing user account:

```
set system login username disable
```

Removing a User Account

From the **su** prompt, use the following command to remove an existing user account:

```
clear system login username
```

Note: The default **admin** account cannot be deleted or disabled.

Changing a Password as an Administrator

From the **su** prompt, use the following command to change the password of any account on the system:

```
set password username
```

Username specifies any account created using the **set system login** command.

Example

This example shows how an administrator would change the password assigned to the "guest" user account:

```
Matrix(su)->set password guest
Please enter new password: ********
Please re-enter new password: ********
Password changed.
Matrix(su)->
```

Changing a Password as an Individual User

From the **rw** prompt, any user with Read-Write access can use the following command to change his or her password:

```
set password
```

Example

This example shows how an individual user would change his or her password:

```
Matrix(rw)->set password
Please enter new password: ********
Please re-enter new password: ********
Password changed.
Matrix(rw)->
```

Verifying Account Information

From the **su** prompt, use the following command to display user account information:

```
show system login
```

Example

This example shows how to display login account information. In this case, device defaults have not been changed:

```
Matrix(su)->show system login
Password history size: 0
Password aging       : disabled

Username        Access      State

admin           super-user  enabled
ro              read-only   enabled
rw              read-write  enabled
```

Setting Password Policies

Once a password is established, the default password policies listed in Table 3-2 apply unless configured otherwise.

Table 3-2 Default Password Policies

Policy	Device Default
Password aging	Disabled
History size (number of passwords checked for duplication)	Disabled
Lock out after failed login attempts	Set to disable Read-Write and Read-Only users — and to lockout the default Admin (Super User) account for 15 minutes — after 3 failed login attempts

Administrators with Super-User access can use the following commands to set password age out time, the password history size, and the number of failed login attempts allowed before user lock out.

Setting Age Out Time

From the **su** prompt, use the following command to enable password aging and set the number of days user passwords will remain valid before aging out, or to disable password aging.

set system password aging {*days* | **disable**}

Valid *days* values are **1** to **365**.

Setting History Size

From the **su** prompt, use the following command to set the number of previously used user login passwords that will be checked for password duplication. This prevents duplicate passwords from being entered into the system with the **set password** command.

set system password history *size*

Valid *size* values are **0** to **10**.

Setting Lockout Attempts and Duration

From the **su** prompt, use the following command to reset the number of failed login attempts allowed and the duration of the timeout. This setting will:

- Disable a Read-Write or Read-Only user account, or
- Lockout the default Admin (**su**) account for a specified number of minutes after maximum login attempts.

  set system lockout {[**attempts** *attempts*] [**time** *time*]}

  Valid *attempts* values are **1** to **10**. Default is **3**.

  Valid *time* values are **0** to **60** (minutes). Default is **15**.

Note: Once a user account is locked out, it can only be re-enabled by a Super User with the **set system login** command (page 3-4).

Starting and Configuring Telnet

In addition to the local console session, the Matrix DFE Series or N-SA device allows up to four simultaneous inbound and/or outbound Telnet sessions.

By default, both inbound and outbound Telnet service is enabled on the device. Use the commands in this section to disable or re-enable Telnet, or to establish a Telnet connection to a remote host.

Enabling or Disabling Telnet

By default, both inbound and outbound Telnet service is enabled on the device. Inbound service sets the ability to Telnet to this device. Outbound service sets the ability to Telnet to other devices. Use the following command to disable or re-enable Telnet:

set telnet {**enable** | **disable**}{**inbound** | **outbound** | **all**}

Disabling Telnet will disconnect all active sessions.

Connecting to a Remote Host Using Telnet

Use the following command to start a Telnet connection from the device's management terminal to a remote host.

telnet *host* [*port*]

If not specified, **23** (the default Telnet *port* number) will be used.

Configuring Secure Shell (SSH) Server

Understanding the SSHv2 Protocol

Secure Shell (SSH) is a "secure" replacement for Telnet. When using Telnet, all communications, including passwords, are sent across the network in clear text (that is, un-encrypted), making eavesdropping on communications an easy task for a knowledgeable user with access to the network. SSH provides the same remote access to the Matrix DFE Series or N-SA device that Telnet provides, but does so securely by encrypting all session data, including passwords. SSH also provides the following additional security features:

- **Public-key authentication of the server**. This feature enables the client to validate the server's authenticity, making it difficult for an attacker to masquerade as the server.
- **Digitally signing all packets**. This feature uses cryptographically strong message digests to authenticate all communications, preventing an attacker from successfully intercepting and altering information.

SSH server is disabled on the Matrix DFE Series and N-SA devices by default. In order to run SSH in its default configuration, you must complete the following steps described in this section.

- Generate host keys
- Enable SSH server

About Host Keys

SSH server authenticates itself to the client through a host key. Host keys are asymmetric encryption keys commonly used in what is known as public key cryptography. SSH server uses unique host keys, each consisting of a pair of keys, generated simultaneously. Although the generated keys are related, one cannot be derived from the other. The first key of the generated pair, the public key, can be published freely and is used by SSH clients to securely identify the SSH server. The second key of the generated pair, the secret key, is stored in a safe place and should never be divulged. This key is used by the SSH server to securely identify itself to SSH clients.

The SSH-2 protocol makes two distinct types of host keys available: the Digital Signature Algorithm (DSA) and the Rivest-Shamir-Adleman (RSA) algorithm. Both DSA and RSA are NIST-approved digital signature algorithms.

After verifying server authenticity, the SSH client generates a key to use until it disconnects from the server. Once the client and server have copies of the key, they will use it to encrypt all further communications. In addition to encrypting each packet, both the client and server will stamp each outgoing packet with data that can be used to validate the contents of the packets. This stamp consists of a message authentication code (MAC) created by using a secure message digest algorithm such as SHA-1 or MD5. If the content of the packet changes en-route, MAC authentication will fail.

Generating Host Keys

Before enabling the SSH server, at least one host key must be generated. Use the following command to generate a host key:

set ssh hostkey

Example

This example shows how to generate SSH private and public host keys. By default, bit size for the DSA and RSA key pairs is 1,024, which is considered very secure:

```
Matrix(rw)->set ssh hostkey

Generating 1024-bit dsa key pair

Key generated.
1024-bit dsa
Private key saved to sshdrv:/.ssh2/dsa
Public key saved to sshdrv:/.ssh2/dsa.pub

Generating 1024-bit rsa key pair

Key generated.
1024-bit rsa
Private key saved to sshdrv:/hostkey
Public key saved to sshdrv:/hostkey.pub
```

Enabling SSH Server

After generating host keys, use the following command to enable SSH Server on the device:

set ssh enable

Setting CLI Parameters

Table 3-3 shows the parameters you can set to control the CLI's display and behavior and their associated commands. Valid values, where applicable, are shown after the command syntax.

Table 3-3 CLI Parameters

Task	Command Syntax <valid values>
Modify the command prompt.	**set prompt** *"prompt string"* **Note:** A prompt string containing a space in the text must be enclosed in quotes.
Set the message of the day banner message displayed at session login.	**set banner motd** *message*
Set the number of columns the CLI will display.	**set width** *screenwidth <50-150>*
Set the number of lines the CLI will display.	**set length** *screenlength <0, 5-512>*
Set the time (in minutes) an idle console or Telnet CLI session will remain connected before timing out.	**set logout** *timeout* <0-333,333>

Table 3-3 CLI Parameters (continued)

Task	Command Syntax <valid values>
Enables the CLI command completion function, which allows you to complete a unique CLI command fragment using the keyboard spacebar.	**set cli completion enable** [**default**]

Displaying Scrolling Screens

If the CLI screen length has been set using the **set length** command as described in Table 3-3, CLI output requiring more than one screen will display `--More--` to indicate continuing screens. To display additional screen output:

- Press any key other than ENTER to advance the output one screen at a time.
- Press ENTER to advance the output one line at a time.

Note: Entering a value of **0** for *screenlength* in the **set length** command will disable scrolling screens.

Example

This example shows how the **show mac** command indicates that output continues on more than one screen.

```
Matrix(rw)->show mac

MAC Address           FID       Port      Type
 ------------------------------------------------------------
 00-00-1d-67-68-69      1       host.0.1  learned
 00-00-02-00-00-00      1       ge.0.2    learned
 00-00-02-00-00-01      1       ge.0.2    learned
 00-00-02-00-00-02      1       ge.0.2    learned
 00-00-02-00-00-03      1       ge.0.2    learned
 00-00-02-00-00-04      1       ge.0.2    learned
 00-00-02-00-00-05      1       ge.0.2    learned
 00-00-02-00-00-06      1       ge.0.2    learned
 00-00-02-00-00-07      1       ge.0.2    learned
 00-00-02-00-00-08      1       ge.0.2    learned
 00-00-02-00-00-09      1       ge.0.2    learned
 00-00-02-00-00-0a      1       ge.0.2    learned
 00-00-02-00-00-0b      1       ge.0.2    learned
--More--
```

Getting Help with CLI Commands

The Matrix DFE Series and N-SA CLI allows you to use context-sensitive help or perform keyword lookups to get help with command sets and syntax.

Using Context-Sensitive Help

Entering **help** after a specific command will display usage and syntax information for that command.

Example

This example shows how to display context-sensitive help for the **set length** command:

```
Matrix(rw)->set length help
Command: set length Number of lines
Usage:  set length  <screenlength>
        screenlength     Length of the screen (5..512, 0 to disable 'more')
```

Performing Keyword Lookups

Entering a space and a question mark (**?**) after a keyword will display all commands beginning with the keyword.

Example

This example shows how to perform a keyword lookup for the **show snmp** command. In this case, 13 additional keywords are used by the **show snmp** command. Entering a space and a question mark (**?**) after any of these parameters (such as **show snmp user**) will display additional parameters nested within the syntax.

```
Matrix(rw)->show snmp ?
  access            SNMP VACM access configuration
  community         SNMP v1/v2c community name configuration
  context           SNMP VACM context list
  counters          SNMP counters
  engineid          SNMP engine properties
  group             SNMP VACM security to group configuration
  notify            SNMP notify configuration
  notifyfilter      SNMP notify filter configuration
  notifyprofile     SNMP notify profile configuration
  targetaddr        SNMP target address configuration
  targetparams      SNMP target parameters configuration
  user              SNMP USM user configuration
  view              SNMP VACM view tree configuration
Matrix(rw)->show snmp
Matrix(rw)->show snmp user ?
  list              List usernames
  <user>            User name
  remote            Show users with remote SNMP engine ID
  volatile          Show temporary entries
  nonvolatile       Show permanent entries
  read-only         Show r/o entries
  <cr>
Matrix(rw)->show snmp user
```

Entering a question mark (**?**) without a space after a partial keyword will display a list of commands that begin with the partial keyword.

Example

This example shows how to use the partial keyword function for all commands beginning with **co**:

```
Matrix(rw)->co?
  configure         Execute a configuration file
  copy              Upload or download an image or configuration file
Matrix(rw)->co
```

Note: At the end of the lookup display, the system will repeat the command you entered without the ?.

Abbreviating and Completing Commands

The Matrix DFE Series and N-SA CLI allows you to abbreviate commands and keywords down to the number of characters that will allow for a unique abbreviation. When the spacebar auto complete function is enabled, it also allows you to determine if a command fragment is unique. If it is, the CLI will complete the fragment on the current display line.

Abbreviating a Command

This example shows how to abbreviate the **show netstat** command to **sh net**.

```
Matrix(rw)->sh net
Active Internet connections (including servers)
Proto Recv-Q Send-Q  Local Address          Foreign Address        State
----- ------ ------  --------------------- ---------------------- -------
TCP        0      0  10.21.73.13.23        134.141.190.94.51246   ESTABLISHED
TCP        0    275  10.21.73.13.23        134.141.192.119.4724   ESTABLISHED
TCP        0      0  *.80                  *.*                    LISTEN
TCP        0      0  *.23                  *.*                    LISTEN
UDP        0      0  10.21.73.13.1030      134.141.89.113.514
UDP        0      0  *.161                 *.*
UDP        0      0  *.1025                *.*
UDP        0      0  *.123                 *.*
```

Using the Spacebar Auto Complete Function

When the spacebar auto complete function is enabled as described in Table 3-3, pressing the spacebar after a CLI command fragment will allow you to determine if the fragment is unique. If it is, the CLI will complete the fragment on the current display line.

This example shows how, when the function is enabled, entering **conf** and pressing the spacebar would be completed as “configure”:

```
Matrix(rw)->conf<SPACEBAR>
Matrix(rw)->configure
```

Using Basic Line Editing Commands

The CLI supports EMACs-like line editing commands. Table 3-4 lists some commonly used line editing keystrokes and their associated commands.

Table 3-4 Basic Line Editing Commands

Key Sequence	Command
Ctrl+A	Move cursor to beginning of line.
Ctrl+B	Move cursor back one character.
Ctrl+C	Abort command.
Ctrl+D	Delete a character.
Ctrl+E	Move cursor to end of line.
Ctrl+F	Move cursor forward one character.
Ctrl+H	Delete character to left of cursor.
Ctrl+I or TAB	Complete word.
Ctrl+K	Delete all characters after cursor.
Ctrl+L or Ctrl+R	Re-display line.
Ctrl+N	Scroll to next command in command history (use the CLI **show history** command to display the history).
Ctrl+P	Scroll to previous command in command history.
Ctrl+T	Transpose characters.
Ctrl+U	Erase entire line.
Ctrl+W	Delete word to the left of cursor.
Ctrl+X	Delete all characters before the cursor.
Ctrl+Y	Restore the most recently deleted item.

Understanding Module and Port Numbering

Important Notice

CLI examples in this guide illustrate a generic Matrix command prompt and chassis-based / modular port designations. Depending on which Matrix device you are using, your default command prompt and output may be different from the examples shown.

Console Ports

Each Matrix Series module or standalone device includes a console port through which local management of the device can be accessed using a terminal or modem.

For module placement rules and considerations for configuring local management on Matrix DFE-Gold Series modules, refer to the *Matrix DFE-Gold Series Installation Guide*.

Switch Ports

Matrix devices have fixed front panel switch ports and, depending on the model, optional expansion module slots. The numbering scheme used to identify the switch ports on the front panel and the expansion module(s) installed is interface-type dependent and, in chassis-based systems, is also dependent upon the chassis in which the module(s) are installed.

Designating Ports in the CLI

Commands requiring a *port-string* parameter use the following syntax to designate port type, location, and port number.

Port String Syntax for Matrix DFE Series Modules

Use this syntax to specify a port string:

```
port type.slot location.port number
```

Where **port type** can be:

- **fe** for 100-Mbps Ethernet
- **ge** for 1-Gbps Ethernet
- **tg** for 10-Gbps Ethernet (Matrix DFE-Platinum 7K4290-02 modules only)
- **com** for COM (console) port(s)
- **host** for host port
- **vlan** for vlan interfaces
- **lag** for IEEE802.3 link aggregation ports
- **lpbk** for loopback interfaces, or
- **lo** for the software loopback interface

Slot location for modules installed in a Matrix N7 or E7 chassis can be:
0 through **7**, with **0** designating virtual system ports (lag, vlan, host), and **1** designating the left-most module slot in the chassis.

Slot location for modules installed in a Matrix N3 chassis can be:
0 through **3**, with **0** designating virtual system ports (lag, vlan, host), and **1** designating the lowest module slot in the chassis.

Port number can be:
Any port number in a slot location. The highest valid port number is dependent on the number of ports in a slot location and the port type.

For example:
If a module in slot 1 has 48, 100 Mbps Fast Ethernet front panel ports, and an uplink interface with 6 Mini-GBICs, the range of port number designations used in the CLI command would be:

- **fe.1.1** through **fe.1.48** for the 48 100 Mbps Fast Ethernet front panel ports, and
- **ge.1.1** through **ge.1.6** for the 6 Mini-GBIC uplink ports.

If the uplink has the same type (**fe**) ports as the front panel, the numbering continues with the port number **fe.1.49**.

Port String Syntax for the Matrix N-SA Device

Use this syntax to specify a port string:

```
port type.port group.port number
```

Where **port type** can be:

- **fe** for 100-Mbps Ethernet
- **ge** for 1-Gbps Ethernet
- **com** for the COM (console) port
- **host** for host port
- **vlan** for vlan interfaces
- **lag** for IEEE802.3 link aggregation ports
- **lpbk** for loopback interfaces, or
- **lo** for the software loopback interface

Port group can be:

- **1** for the lower fixed front panel ports
- **2** for the middle fixed front panel ports, or
- **3** for the top fixed front panel ports, and the Mini-GBIC uplink ports

For example:
If a module in slot 1 has 48, 100 Mbps Fast Ethernet front panel ports, and an uplink interface with four Mini-GBICs, the range of port number designations used in the CLI command would be:

- **fe.1.1** through **fe.1.20** for the lower fixed front panel Fast Ethernet ports
- **fe.2.1** through **fe.2.20** for the middle fixed front panel Fast Ethernet ports
- **fe.3.1** through **fe.3.8** for the top fixed front panel Fast Ethernet ports, and
- **ge.3.14**, **ge.3.16**, **ge.3.17**, or **ge.3.17** for the Mini-GBIC uplink ports.

Note: You can use a wildcard (*) to indicate all of an item. For example, fe.3.* would represent all 100 Mbps Ethernet (fe) ports on the module in slot/port group 3.

Table 3-5 shows the *port-string* syntax for specifying various ports on Matrix devices. Your options will vary, as described above, depending on your device(s) and system configuration:

Table 3-5 Examples of Port String Designations

Port Type and Location	Port-String Syntax
Fast Ethernet ports 1 through 10 in slot/port group 1	`fe.1.1-10`
1-Gigabit Ethernet port 14 in slot /port group 3	`ge.3.2`
10-Gigabit Ethernet port 2 of the module in slot 3	`tg.3.1`
Fast Ethernet ports 1, 3, 7, 8, 9 and 10 in the module in chassis slot 1	`fe.1.1,fe.1.3,fe.1.7-10`
All 1-Gigabit Ethernet ports slot/port group 3	`ge.3.*`
All 10-Gigabit Ethernet ports in the chassis	`tg.*.*`

Table 3-5 Examples of Port String Designations (continued)

Port Type and Location	Port-String Syntax
All ports (of any interface type) of all modules in the chassis or standalone device	`*.*.*`
The console port in module 1 or in a standalone device	`com.1.1`
Virtual LAG port 2	`lag.0.2`

Preparing the Device for Router Mode

Pre-Routing Configuration Tasks

The following pre-routing tasks must be performed from the switch CLI.

- Starting up the CLI. (page 2-1)
- Setting the system password. (page 2-4)
- Configuring basic platform settings, such as host name, system clock, and terminal display settings. (page 2-4)
- Setting the system IP address. (page 2-4)
- Creating and enabling VLANs.
- File management tasks, including uploading or downloading flash or text configuration files, and displaying directory and file contents. (page 2-7)
- Configuring at least one module to run in router mode. (page 3-16)

Notes: The command prompts used as examples in this section and throughout this guide show switch operation for a user in Read-Write (rw) access mode, and a system where module 1 (or the N-SA device) and VLAN 1 have been configured for routing. The prompt changes depending on your current configuration mode, the specific module, and the interface types and numbers configured for routing on your system.

A *module* designation of 1 must be entered to enable routing on the Matrix N-SA standalone device. All other values will result in an error message.

Use the following procedure and associated commands to enable the switch for routing.

Procedure 3-1 Enabling the Switch for Routing

Step	Task	Command(s)
1.	**In switch mode:** Configure a routing module. **Note:** A module designation of 1 must be entered to enable routing on the Matrix N-SA standalone device.	**`set router`** *`module`*
2.	Enable router mode.	**`router`** *`module`*

Procedure 3-1 Enabling the Switch for Routing (continued)

Step	Task	Command(s)
3.	**In router mode:** Enable router Privileged EXEC mode.	**enable**
4.	**In router Privileged EXEC mode**: Enable global router configuration mode.	**configure terminal**
5.	**In Global Configuration mode:** Enable interface configuration mode using the interface of the routing module.	**interface** {**vlan** *vlan-id* | **loopback** *loopback-id*
6.	**In Interface Configuration mode:** Assign an IP address to the routing interface.	**ip address** {*ip-address ip-mask*}
7.	Enable the interface for IP routing.	**no shutdown**

Example

This example shows how to:

- Configure Matrix DFE Series module 1, or the N-SA device, as a routing module.
- Configure VLAN 1 on IP address 182.127.63.1 255.255.255.0 as the routing interface for that module.

```
Matrix(rw)->set router 1
Matrix(rw)->router 1
Matrix>Router1>enable
Matrix>Router1#configure terminal
 Enter configuration commands:
Matrix>Router1(config)#interface vlan 1
Matrix>Router1(config-if(Vlan 1))#ip address 182.127.63.1 255.255.255.0
Matrix>Router1(config-if(Vlan 1))#no shutdown
```

Enabling Router Configuration Modes

The Matrix DFE / N-SA Series CLI provides different modes of router operation for issuing a subset of commands from each mode. Table 3-6 describes these modes of operation.

Table 3-6 Router Configuration Modes

Use this mode...	To...	Access method...	Resulting Prompt...
Privileged EXEC Mode	Set system operating parameters Show configuration parameters Save/copy configurations	From the switch CLI: 1. Type **router** *module* (using a module number configured for routing). 2. Type **enable**.	`Matrix>Router1>` `Matrix>Router1#`

Table 3-6 Router Configuration Modes (continued)

Use this mode...	To...	Access method...	Resulting Prompt...
Global Configuration Mode	Set system-wide parameters.	Type **configure terminal** from Privileged EXEC mode.	`Matrix>Router1(config)#`
Interface Configuration Mode	Configure router interfaces.	Type **interface vlan** or **interface loopback** and the interface's *id* from Global Configuration mode.	`Matrix>Router1 (config-if(Vlan 1 | Lpbk 1))#`
Router Configuration Mode	Set IP protocol parameters.	Type **router** and the *protocol name* from Global or Interface Configuration mode.	`Matrix>Router1 (config-router)#`
Key Chain Configuration Mode	Set protocol (RIP) authentication key parameters.	Type **key chain** and the key chain *name* from Router (RIP) Configuration mode.	`Matrix>Router1 (config-keychain)#`
Key Chain Key Configuration Mode	Configure a specific key within a RIP authentication key chain.	Type **key** and the *key-id* from Key Chain Configuration Mode.	`Matrix>Router1 (config-keychain-key)#`
Policy-Based Routing Configuration Mode	Configure policy-based routing for a route map.	Type **route-map**, an *id-number*, and **permit** or **deny** from Global Configuration Mode.	`Matrix>Router1 (config-route-map-pbr)#`
Server Load Balancing (SLB) Server Farm Configuration Mode	Configure an LSNAT server farm.	Type **ip slb serverfarm** and the *serverfarmname* from Global Configuration Mode.	`Matrix>Router1 (config-slb-sfarm)#`
Server Load Balancing (SLB) Real Server Configuration Mode	Configure an LSNAT real server.	Type **real** and the real server *IP address* from SLB Server Farm Configuration Mode.	`Matrix>Router1 (config-slb-real)#`
Server Load Balancing (SLB) Virtual Server Configuration Mode	Configure an LSNAT virtual server.	Type **ip slb vserver** and the *vserver-name* from Global Configuration Mode.	`Matrix>Router1 (config-slb-vserver)#`
IP Local Pool Configuration Mode	Configure a local address pool as a DHCP subnet	Type **ip local pool** and the local pool *name* from Global Configuration Mode.	`Matrix>Router1 (ip-local-pool)#`

Table 3-6 Router Configuration Modes (continued)

Use this mode...	To...	Access method...	Resulting Prompt...
DHCP Pool Configuration Mode	Configure a DHCP server address pool.	Type **ip dhcp pool** and the address pool *name* from Global Configuration Mode.	`Matrix>Router1 (config-dhcp-pool)#`
DHCP Class Configuration Mode	Configure a DHCP client class.	Type **client-class** and the client class *name* from DHCP Pool or Host Configuration Mode.	`Matrix>Router1 (config-dhcp-class)#`
DHCP Host Configuration Mode	Configure DHCP host parameters.	Type **client-identifier** and the *identifier*, or **hardware-address** and an *address* from any DHCP configuration mode.	`Matrix>Router1 (config-dhcp-host)#`

Note: To jump to a lower configuration mode, type **exit** at the command prompt. To revert back to switch CLI, type **exit** from Privileged EXEC router mode.

4

Configuring Link Aggregation

This chapter provides information about the following link aggregation configuration procedures on the Matrix DFE Series or N-SA device.

For information about...	Refer to page...
Link Aggregation Control Protocol (LACP) Overview	4-1
Configuring LAG Aggregator Keys and Priority	4-4
Configuring Underlying Physical Ports	4-5
Enabling LACP Flow Regeneration	4-7

Link Aggregation Control Protocol (LACP) Overview

Caution: Link aggregation configuration should only be performed by personnel who are knowledgeable about Spanning Tree and Link Aggregation, and fully understand the ramifications of modifications beyond device defaults. Otherwise, the proper operation of the network could be at risk.

Using multiple links simultaneously to increase bandwidth is a desirable switch feature, which can be accomplished if both sides agree on a set of ports that are being used as a Link Aggregation Group (LAG). Once a LAG is formed from selected ports, problems with looping can be avoided since the Spanning Tree can treat this LAG as a single port.

The Link Aggregation Control Protocol (LACP) logically groups interfaces together to create a greater bandwidth uplink, or link aggregation, according to the IEEE 802.3ad standard. This standard allows the switch to determine which ports are in LAGs and configure them dynamically. Since the protocol is based on the IEEE 802.3ad specification, any switch from any vendor that supports this standard can aggregate links automatically.

802.3ad LACP aggregations can also be run to end-users (for example, a server) or to a router.

Note: Earlier (proprietary) implementations of port aggregation referred to groups of aggregated ports as “trunks.”

LACP Operation

For each aggregatable port in the device, LACP:

- Maintains configuration information (reflecting the inherent properties of the individual links as well as those established by management) to control aggregation.
- Exchanges configuration information with other devices to allocate the link to a Link Aggregation Group (LAG).

Note: A given link is allocated to, at most, one Link Aggregation Group (LAG) at a time. The allocation mechanism attempts to maximize aggregation, subject to management controls.

- Attaches the port to the aggregator used by the LAG, and detaches the port from the aggregator when it is no longer used by the LAG.
- Uses information from the partner device's link aggregation control entity to decide whether to aggregate ports.

The operation of LACP involves the following activities:

- Checking that candidate links can actually be aggregated.
- Controlling the addition of a link to a LAG, and the creation of the group if necessary.
- Monitoring the status of aggregated links to ensure that the aggregation is still valid.
- Removing a link from a LAG if its membership is no longer valid, and removing the group if it no longer has any member links.

In order to allow LACP to determine whether a set of links connect to the same device, and to determine whether those links are compatible from the point of view of aggregation, it is necessary to be able to establish the following:

- A globally unique identifier for each device that participates in link aggregation
- A means of identifying the set of capabilities associated with each port and with each aggregator, as understood by a given device
- A means of identifying a LAG and its associated aggregator

LACP Terminology

Table 4-1 defines key terminology used in LACP configuration.

Table 4-1 LACP Terms and Definitions

Term	Definition
Aggregator	Virtual port that controls link aggregation for underlying physical ports. Each Matrix DFE-Platinum Series module or N-SA device provides 48 aggregator ports, which are designated in the CLI as **lag.0.1** through **lag.0.48**. Each Matrix DFE-Gold Series module provides 4 aggregator ports, which are designated in the CLI as **lag.0.1** through **lag.0.4**.
LAG	Link Aggregation Group. Once underlying physical ports (such as, **fe.x.x**, or **ge.x.x**) are associated with an aggregator port, the resulting aggregation will be represented as one LAG with a **lag.x.x** port designation.

Table 4-1 LACP Terms and Definitions (continued)

Term	Definition
LACPDU	Link Aggregation Control Protocol Data Unit. The protocol exchanges aggregation state/mode information by way of a port's actor and partner operational states. LACPDUs sent by the first party (the actor) convey to the second party (the actor's protocol partner) what the actor knows, both about its own state and that of its partner.
Actor and Partner	An actor is the local device sending LACPDUs. Its protocol partner is the device on the other end of the link aggregation. Each maintains current status of the other using LACPDUs containing information about their ports' LACP status and operational state.
Admin Key	Value assigned to aggregator ports and physical ports that are candidates for joining a LAG. The LACP implementation on Matrix DFE Series and N-SA devices will use this value to form an operational key and will determine which underlying physical ports are capable of aggregating by comparing operational keys. Aggregator ports allow only underlying ports with operational keys matching theirs to join their LAG.
System Priority	Value used to build a LAG ID, which determines aggregation precedence. If there are two partner devices competing for the same aggregator, LACP compares the LAG IDs for each grouping of ports. The LAG with the lower LAG ID is given precedence and will be allowed to use the aggregator. **Note:** Only one LACP system priority can be set on a Matrix DFE Series or N-SA device, using either the **set lacp asyspri** command (page 4-6), or the **set port lacp** command (page 4-6).
Port Priority	Used by the device with the lowest LAG ID to determine which underlying physical ports will be allowed into a LAG. Ports with the lowest priority are allowed to join and all others are placed in standby.

Matrix DFE Series and N-SA Usage Considerations

In normal usage (and typical implementations) there is no need to modify any of the default LACP parameters on the Matrix DFE Series and N-SA device. The default configurations, as shown in Table 4-2, will result in the maximum number of aggregations possible. If the switch is placed in a configuration with peers not running the protocol, no dynamic link aggregations will be formed and the switch will function normally (that is, will block redundant paths) .

Table 4-2 Link Aggregation Configurations by Device

Device	Aggregator Ports Provided	Ports Allowed in a LAG
Matrix DFE-Platinum Series	48 (**lag.0.1** - **lag.0.48**)	unlimited
Matrix N-SA	48 (**lag.0.1** - **lag.0.48**)	unlimited
Matrix DFE-Gold Series	4 (**lag.0.1** - **lag.0.4**)	4

As shown in Table 4-2, depending on the model, each Matrix device provides either 4 or 48 aggregator ports, which are designated in the CLI as **lag.0.1** through **lag.0.48**. Once underlying physical ports (such as, **fe.x.x**, or **ge.x.x**) are associated with an aggregator port, the resulting aggregation will be represented as one LAG with a **lag.x.x** port designation. LACP determines which underlying physical ports are capable of aggregating by comparing operational keys. Aggregator ports allow only underlying ports with keys matching theirs to join their LAG.

LACP uses a system priority value to build a LAG ID, which determines aggregation precedence. If there are two partner devices competing for the same aggregator, LACP compares the LAG IDs for each grouping of ports. The LAG with the lower LAG ID is given precedence and will be allowed to use the aggregator.

There are a few cases in which ports will not aggregate:

- An underlying physical port is attached to another port on this same switch (loopback).
- There is no available aggregator for two or more ports with the same LAG ID. This can happen if there are simply no available aggregators, or if none of the aggregators have a matching admin key and system priority.
- 802.1x authentication is enabled using the **set eapol** command and ports that would otherwise aggregate are not 802.1X authorized.

As shown in Table 4-2, the LACP implementation on Matrix DFE-Platinum Series modules or the N-SA device will allow as many ports as possible into a LAG, while the implementation on Matrix DFE-Gold Series modules will allow up to a maximum of four ports into a LAG. The device with the lowest LAG ID determines which underlying physical ports are allowed into a LAG based on the ports' LAG port priority. Ports with the lowest LAG port priority values are allowed into the LAG and all others go into a standby state.

Note: To aggregate, underlying physical ports must be running in full duplex mode and must be of the same operating speed.

Configuring LAG Aggregator Keys and Priority

Reviewing the Configuration

By default, LACP will be enabled on all ports. Aggregator admin key and system priority values will both be set at 32768, and no LAGs will be configured.

Use this command to review the current configuration:

show lacp [*port-string*]

Example

This example shows how to display information for aggregator port 4. In this case, default LACP parameters have not been changed:

```
Matrix(rw)->show lacp lag.0.4

Aggregator:  lag.0.4
                          Actor                 Partner
System Identifier:  00:e0:63:9d:b5:87      00:00:00:00:00:00
  System Priority:              32768                  32768
        Admin Key:              32768
         Oper Key:              32768                  32768
   Attached Ports:  None.
```

Assigning a LAG Aggregator Key

LACP allows only underlying physical ports with keys that match their aggregators to join a LAG. You can change the default aggregator admin key for one or more ports, allowing those ports to join a LAG with a matching key value.

Use this command to change the aggregator key for one or more LAG ports:

set lacp aadminkey lag*port-string value*

The *lagport-string* must specify a virtual LAG port (lag.0.*x*)

Valid aggregator *values* are 1 to 65535, with precedence given to lower values.

Assigning a LAG System Priority

LACP uses system priority to build a LAG ID which determines aggregation precedence. When two partner devices compete for the same aggregator, LACP compares the LAG IDs for each grouping of ports. The LAG with the lower ID is given precedence, and will be allowed to use the aggregator.

Use this command to change the priority for an aggregator:

set lacp asyspri *value*

Valid *values* are 0 to 65535, with precedence given to lower values.

Note: Only one LACP system priority can be set on a Matrix DFE Series or N-SA device, using either the **set lacp asyspri** or the **set port lacp** commands.

Configuring Underlying Physical Ports

In normal usage (and typical implementations) there is no need to modify any of the default LACP parameters on the Matrix DFE Series or N-SA device. The default values will result in the maximum number of aggregations possible. If you are knowledgeable about link aggregation configuration and want to control LACP behavior through the CLI, you can use the commands in this section to create a static LAG, or to modify parameters for one or more underlying physical ports.

Statically Assigning Ports to a LAG

Use this command to assign one or more underlying physical ports to a LAG:

set lacp static *lagportstring* [**key** *key*] *port-string*

The *key* value specifies the new member port and LAG port aggregator admin key value. This must be a unique value between 0 and 65535. If ports other than the desired underlying physical ports share the same admin key value, aggregation will fail, or undesired aggregations will form.

Notes: At least two ports need to be assigned to a LAG port for a Link Aggregation Group to form and attach to the specified LAG port. Matrix DFE-Gold Series devices allow a maximum of four ports per aggregator.

The same usage considerations for dynamic LAGs discussed on page 4-3 apply to statically created LAGs.

Example

This example shows how to add Fast Ethernet port 6 in slot 1 to the LAG of aggregator port 4:

```
Matrix(rw)->set lacp static lag.0.4 fe.1.6
```

Setting Underlying LACP Port Parameters

Link aggregation parameters can be set for underlying physical ports that will determine their ability to join a LAG, and their administrative state once aggregated.

Use the following commands to set link aggregation actor and partner parameters for one or more underlying ports.

About actor and partner settings

LACP commands and parameters beginning with an "a" (such as **aadminkey**) set actor values. Corresponding commands and parameters beginning with a "p" (such as **padminkey**) set corresponding partner values.

Actor refers to the local device participating in LACP negotiation, while partner refers to its remote device partner at the other end of the negotiation. Actors and partners maintain current status of the other using LACPDUs containing information about their ports' LACP status and operational state.

To set the actor aggregator key:

set port lacp port port-string aadminkey *aadminkey*

The *aadminkey* value (0 - 65535) must match that set for a LAG port in order for the specified port(s) to join the LAG.

To set the system or port priorities:

set port lacp port *port-string* [**asyspri** *asyspri*][**aportpri** *aportpri*] [**padminportpri** *padminportpri*]}

The *asyspri* value (0 - 65535) must match the *aadminkey* value set for a LAG port in order for the specified port(s) to join the LAG.

The *aportpri* and *padminportpri* (actor and partner priority) values (0 - 65535) will be used by the device with the lowest system priority to determine which ports will be given precedence to join a LAG.

Note: Only one LACP system priority can be set on a Matrix DFE Series or N-SA device, using either the **set lacp asyspri** or the **set port lacp** commands.

To set the actor or partner administrative states:

```
set port lacp port port-string aadminstate | padminstate {lacpactive |
lacptimeout | lacpagg | lacpsync | lacpcollect | lacpdist | lacpdef |
lacpexpire}
```

Admin state parameters allow for actor or partner ports to perform as follows:

- **lacpactive** - Transmit LACP PDUs.
- **lacptimeout** - Transmit LACP PDUs every 1 second versus 30 seconds (default).
- **lacpagg** - Aggregate on this port.
- **lacpsync** - Transition to synchronization state.
- **lacpcollect** - Transition to collection state.
- **lacpdist** - Transition to distribution state.
- **lacpdef** - Transition to defaulted state.
- **lacpexpire** - Transition to expired state.

Enabling LACP Flow Regeneration

The flow regeneration function on the Matrix DFE Series or N-SA device allows LACP to redistribute all existing flows over a link aggregation group when a new port joins the LAG. Flow regeneration will also attempt to load balance existing flows to take advantage of ports added to the LAG. By default, flow regeneration is disabled on the device. This means that when a new port joins a LAG, LACP will only distribute new flows over the increased number of ports in the LAG and will leave existing flows intact.

Use this command to enable LACP flow regeneration:

```
set lacp flowRegeneration enable
```

5

Configuring Spanning Trees

This chapter provides the following information about configuring and monitoring Spanning Tree protocols on the Matrix DFE Series and N-SA device:

For information about...	Refer to page...
Overview of Spanning Tree Protocols	5-1
Configuring STP and RSTP	5-8
Configuring MSTP	5-14
Configuring Spanguard	5-16

Overview of Spanning Tree Protocols

Matrix DFE Series and N-SA devices support the Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP) as defined in the following standards:

- IEEE 802.1D (Spanning Tree Protocol)
- IEEE 802.1w (Rapid Spanning Tree Protocol)
- IEEE 802.1s (Multiple Spanning Tree Protocol)

Note: MSTP and RSTP are fully compatible and interoperable with each other and with legacy STP 802.1D.

Spanning Tree Terms and Definitions

Table 5-1 lists terms and definitions used in spanning tree configuration.

Table 5-1 Spanning Tree Terms and Definitions

Term	Definition
Alternate port	Acts as an MSTP alternate path to the root bridge than that provided by the root port.
Backup port	Acts as an MSTP backup for the path provided by a designated port toward the leaves of the spanning tree. Backup ports can exist only where two ports are connected together in a loopback mode or bridge with two or more connections to a shared LAN segment.
BID	Bridge identification, which is derived from the bridge's MAC address and bridge priority. The bridge with the lowest BID becomes the root bridge.

Table 5-1 Spanning Tree Terms and Definitions (continued)

Term	Definition
BPDU	Bridge Protocol Data Unit messages. Used by STP to exchange information, including designating a bridge for each switched LAN segment, and one root bridge for the spanning tree.
Bridge	Switching device.
Bridge priority	Assigns the bridge's relative priority compared to other bridges.
CIST	Common and Internal Spanning Tree created by MSTP to represent the connectivity of the entire network. This is equivalent to the single spanning tree used for STP and RSTP. Communications between MST regions occurs using the CIST.
Designated port	A forwarding port within an active topology elected for every switched LAN segment.
Edge port	Port on the edge of a bridged LAN.
FID	Filter Identifier. Each VLAN is associated to a FID. VLANs are mapped to SIDs using their FID association.
Forward delay	Time interval (in seconds) the bridge spends in listening or learning mode before it begins forwarding BPDUs.
Hello time	Time interval (in seconds) at which the bridge sends BPDUs.
ISLs	Inter-Switch Links.
Master port	The root port for an entire MST region, providing connectivity from the MST region to a CIST root that lies outside the MST region.
Max age	Maximum time (in seconds) the bridge can wait without receiving a configuration message (bridge "hello") before attempting to reconfigure.
MST region	An MSTP group of devices configured together to form a logical region. The MST region presents itself to the rest of the network as a single device, which simplifies administration.
MSTI	Multiple Spanning Tree Instance. Matrix DFE Series and N-SA devices support up to 64 MSTIs.
Path cost	Sum of the port costs in a path to a specified destination.
Port cost	Value assigned to a port based on the speed of the port. The faster the speed, the lower the cost. This helps to determine the quickest path between the root bridge and a specified destination. The segment attached to the root bridge normally has a path cost of zero.
Port priority	Assigns a port's priority in relation to the other ports on the same bridge.
Root bridge	Logical center of the spanning tree, used by STP to determine which paths to block and which to open.
Root port	Port in an active topology through which the root bridge can be reached.
SID	Spanning tree identifier. By default, SID 0 is assumed. VLANs are mapped to SIDs using their FID association.

The following sections describe the spanning tree protocols supported by the Matrix DFE Series or N-SA device.

Spanning Tree (IEEE 802.1D)

The Spanning Tree Protocol (STP) defined in IEEE 802.1D allows bridges to dynamically discover a subset of the topology that is loop-free. The loop-free tree that is discovered contains paths to every LAN segment.

The Spanning Tree Protocol is used to eliminate data loops in an Ethernet network by creating a tree where there is only one data route between any two end stations. STP blocks redundant data paths. Should a path become unreachable, STP automatically activates a blocked path. Should a bridge be added creating a redundant path, STP blocks one of the paths. STP can also change data paths based on a cost change.

All bridges that support the spanning tree exchange information using Bridge Protocol Data Unit (BPDU) messages. Using the information exchanged by the BPDUs, STP designates a bridge for each switched LAN segment, and one root bridge for the spanning tree. The root bridge is the logical center of the spanning tree and is used to determine which paths to block and which to open.

A network administrator can determine the topology of the spanning tree by adjusting the bridge priority, port priority, and path cost. The bridge priority assigns the bridge's relative priority compared to other bridges. The port priority assigns the port's priority in relation to the other ports on the same bridge. By default, the port cost is a value assigned to the port based on the speed of the port. The faster the speed, the lower the cost. This helps to determine the quickest path between the root bridge and a specified destination. The segment attached to the root bridge normally has a path cost of zero.

Note: The terms *path cost* and *port cost* are sometimes used interchangeably. Normally, the port cost is the value assigned to a specific port, while path cost, especially in a BPDU, is the sum of the port costs in a path.

Each bridge has a Bridge Identification (BID), which is derived from the bridge's MAC address and bridge priority. The bridge with the lowest BID becomes the root bridge.

Rapid Spanning Tree (IEEE 802.1w)

The Rapid Spanning Tree Protocol (RSTP) defined in IEEE 802.1w enhances the STP by allowing the network topology to reconverge in a significantly smaller amount of time. This is partially accomplished by its ability to be independent of the protocol timer values when configuring the active topology of a LAN.

Multiple Spanning Trees (IEEE 802.1s)

The Multiple Spanning Tree Protocol (MSTP), defined in IEEE 802.1s, allows for increased bandwidth utilization and optimal load balancing across redundant links. It further expands upon STP and RSTP with the following features:

- Backwards compatibility with STP and RSTP.
- Ability to create a single Common and Internal Spanning Tree (CIST) that represents the connectivity of the entire network.
- Users can group any number of devices into individual regions, with each region behaving and presenting itself as a single device to the rest of the network.
- A region can contain multiple instances of the spanning tree, where each instance can support multiple VLANs.

- In addition to using hello time, forward delay, and max age information, MSTP also utilizes the hop count for improved performance.

MSTP can automatically detect the version of spanning tree being used on a LAN and send out the equivalent type of BPDU. In addition, MSTP incorporates a force version feature where the user may force MSTP to behave as STP or RSTP.

Common and Internal Spanning Tree (CIST)

MSTP uses all STP, RSTP and Multiple Spanning Tree (MST) region information to create a single Common and Internal Spanning Tree (CIST) that represents the connectivity of the entire network. This is equivalent to the single spanning tree used for STP and RSTP.

The MSTP enabled network contains one CIST and a minimum of at least one MST region. A typical network may contain numerous MST regions as well as separate LAN segments running legacy STP and RSTP spanning tree protocols.

The CIST contains a root bridge, which is the root of the spanning tree for the network. The CIST root is not necessarily located inside an MST region. Each region contains a CIST regional root, unless the CIST root is part of the region.

MST Region

An MST region is a group of devices that are configured together to form a logical region. The MST region presents itself to the rest of the network as a single device, which simplifies administration. Path cost is only incremented when traffic enters or leaves the region, regardless of the number of devices within the region. Each LAN can only be a member of one region. Figure 5-1 shows that the MST region appears as a single device to Devices 1 and 2, but really consists of three devices.

Figure 5-1 Example of an MST Region

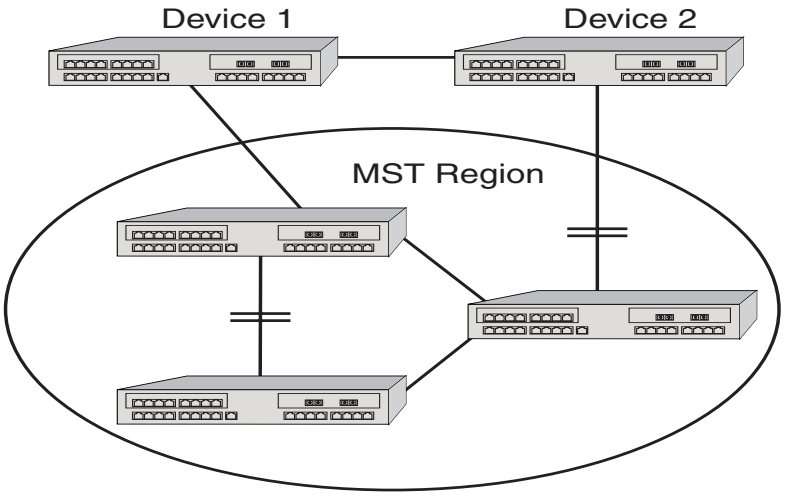

For a device to be considered as part of an MST region, it must be administratively configured with the same configuration identifier information as all other devices in the MST region. The configuration identifier consists of four separate parts:

- Format Selector - One octet in length and is always 0. It cannot be administratively changed.
- Configuration Name - A user-assigned, case sensitive name given to the region. The maximum length of the name is 32 octets.
- Revision Level - Two octets in length. The default value of 0 may be administratively changed.
- Configuration Digest - 16 octet HMAC-MD5 signature created from the configured VLAN Identification (VID)/Filtering Identification (FID) to Multiple Spanning Tree Instances (MSTI) mappings. All devices must have identical mappings to have identical configuration digests.

The MST region designates one CIST regional root bridge for the region, regardless of the number of MSTIs. The regional root provides the connectivity from the region to the CIST root, when the CIST root lies outside the region.

Multiple Spanning Tree Instances (MSTI)

Inside the MST region, a separate topology is maintained from the outside world. Each MST region may contain up to 64 different MSTIs. The Matrix DFE Series or N-SA device maps VLAN IDs (VIDs) and Filtering IDs (FIDs) to each other in a one to one correlation; for example, FID 3 = VID 3. VID/FIDs are mapped to different MSTIs to create a type of load balancing.

Determining FID-to-SID Mappings

VLANs are mapped to MSTIs through a FID-to-SID mapping correlation which is the key element in MSTP configuration. Each VLAN is associated to a FID and, during MSTI creation, VLANs are mapped to spanning tree IDs using their FID association. This mapping is contained within the MST configuration digest described in the previous section and displayed in the following example. By default, every bridge will have a FID-to-SID mapping that equals VLAN FID 1/SID 0.

Use this command to determine MSTI configuration identifier information, and whether or not there is a misconfiguration due to non-matching configuration identifier components:

```
show spantree mstcfgid
```

Example

This example shows how to display MSTI configuration identifier information. In this case, this bridge belongs to “Region1”:

```
Matrix->show spantree mstcfgid
MST Configuration Identifier:
 Format Selector:      0
 Configuration Name:   Region1
 Revision Level:       88
 Configuration Digest: 6d:d7:93:10:91:c9:69:ff:48:f2:ef:bf:cd:8b:cc:de
```

In order for other bridges to belong to Region1, all four elements of those bridges’ configuration ID output must match. The only default value that must be changed for this to happen is the configuration name setting, as follows:

Use this command to change the configuration name from the default bridge MAC address value to “Region 1”:

```
set spantree mstcfgid cfgname Region1
```

Since an MSTI is a separate spanning tree, each MSTI has its own root inside the MST region. Figure 5-2 and Figure 5-3 show two MSTIs in a single region. Device 3 is the root for MSTI 1, Device 2 is the root for MSTI 2, and Device 5 is the CIST regional root. Traffic for all the VLANs attached to an MSTI follow the MSTI’s spanned topology.

Various options may be configured on a per-MSTI basis to allow for differing topologies between MSTIs. To reduce network complexity and processing power needed to maintain MSTIs, you should only create as many MSTIs as needed.

Figure 5-2 MSTI 1 in a Region

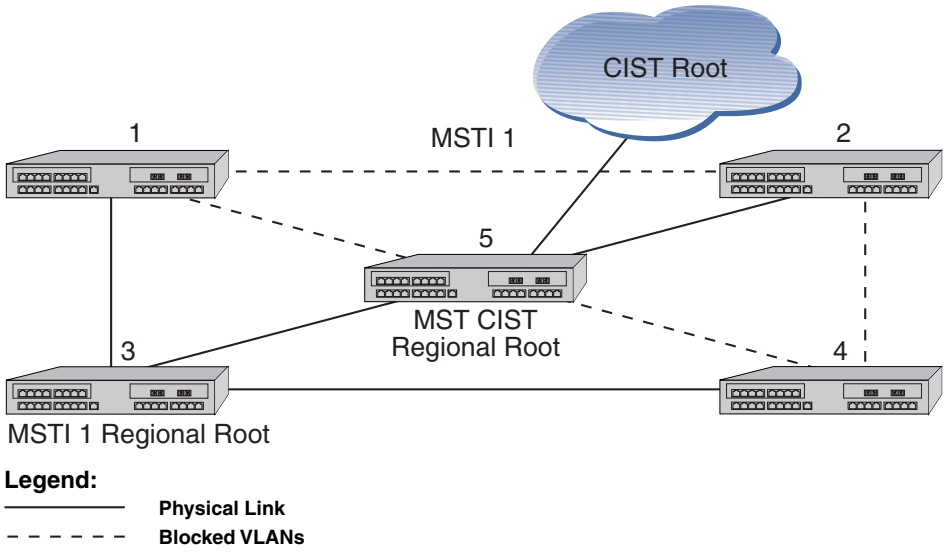

Figure 5-3 MSTI 2 in the Same Region

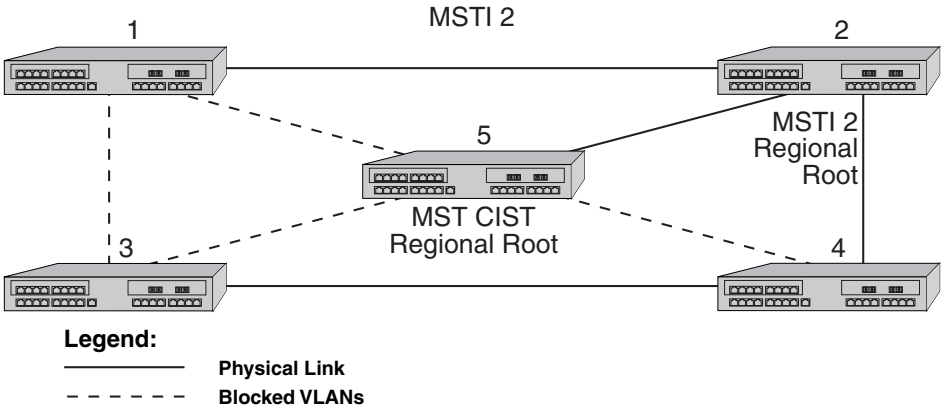

Figure 5-4 shows 3 regions with five MSTIs. Table 5-2 defines the characteristics of each MSTI. Ports connected to PCs from Devices 1, 3, 9, and 11 will be automatically detected as edge ports. Devices 4 and 10 are the CIST regional roots and, because they contain the master port for their regions, are also the regional root devices. Each MSTI can be configured to forward and block various VLANs.

Figure 5-4 Example of Multiple Regions and MSTIs

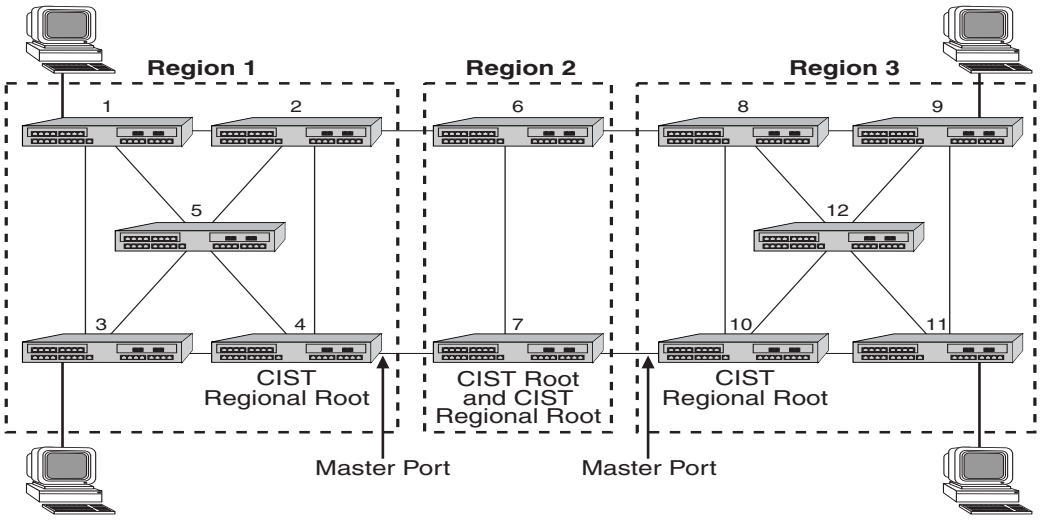

Table 5-2 MSTI Characteristics for Figure 5-4

MSTI / Region	Characteristics
MSTI 1 in Region 1	Root is Device 4, which is also the CIST regional root
MSTI 2 in Region 1	Root is Device 5
MSTI 1 in Region 2	Root is Device 7, which is also the CIST root
MSTI 1 in Region 3	Root is Device 11
MSTI 2 in Region 3	Root is Device 12
	Device 10 is the CIST regional root

Spanning Tree Port States

Spanning tree assigns port states to control the forwarding and learning processes within a topology. Table 5-3 lists Spanning Tree port states, and the behavior associated with each state:

Table 5-3 Spanning Tree Port States

Port State	Behavior
Blocking	Actively preventing traffic from using this path.
Listening	Continuing to block traffic while waiting for protocol information to determine whether to go back to the blocking state, or continue to the learning state.
Learning	Learning station location information but continuing to block traffic.
Forwarding	Forwarding traffic and continuing to learn station location information.
Disabled	Disabled administratively or by failure.
Discarding	MSTP state with the same behavior as blocking.

Configuring STP and RSTP

Caution: Spanning Tree configuration should be performed only by personnel who are very knowledgeable about Spanning Trees and the configuration of the Spanning Tree Algorithms. Otherwise, the proper operation of the network could be at risk.

This section provides information about the following STP and RSTP tasks:

- Reviewing and Re-enabling Spanning Tree
- Adjusting Spanning Tree Parameters
- Enabling the Backup Root Function
- Adjusting RSTP Parameters

Reviewing and Re-enabling Spanning Tree

By default, spanning tree is enabled globally on the Matrix DFE Series or N-SA device and enabled on all ports. The default version is set to MSTP (802.1s) mode. Since MSTP mode is fully compatible and interoperable with legacy STP 802.1D and RSTP bridges, in most networks, this default should not be changed.

Use the following commands to review, re-enable and reset the spanning tree mode.

1. Review the current configuration on one or more SIDs and/or ports:

   show spantree stats [**port** *port-string*][**sid** *sid*][**active**]

   Specifying **active** will display information for port(s) that have received BPDUs since boot.

Example

This example shows how to display the device's spanning tree configuration:

```
Matrix->show spantree stats

SID                          - 1
Spanning tree mode           - enabled
Designated Root              - 00-e0-63-6c-9b-6d
Designated Root Priority     - 0
Designated Root Cost         - 1
Designated Root Port         - ge.5.1
Root Max Age                 - 20 sec
Root Hello Time              - 2  sec
Root Forward Delay           - 15 sec
Bridge ID MAC Address        - 00-e0-63-9d-b5-87
Bridge priority              - 32768
Bridge Max Age               - 20 sec
Bridge Hello Time            - 2  sec
Bridge Forward Delay         - 15 sec
Topology Change Count        - 6539
Time Since Top Change        - 00 days 00:00:00
```

2. If necessary, globally enable spanning tree:

   set spantree stpmode iee8021

3. Review the status of spanning tree on one or more ports:

   show spantree portenable [**port** *port-string*]

4. If necessary, re-enable spanning tree on one or more ports:

   set spantree portenable *port-string* **enable**

Adjusting Spanning Tree Parameters

You may need to adjust certain spanning tree parameters if the default values are not suitable for your bridge configuration. Parameters affecting the entire spanning tree are configured with variations of the global bridge configuration commands. Interface-specific parameters are configured with variations of the spanning tree port configuration commands. Default settings are:

- Bridge priority: 32768
- Port priority: 128
- Port cost - 0 (automatically calculated based on port speed)
- Bridge hello time - 2 seconds
- Bridge forward delay interval - 15 seconds
- Bridge maximum aging time - 20 seconds
- Maximum spanning tree instances allowed:
  - 1 - 9 on Matrix DFE-Gold Series modules;
  - 1 - 33 on Matrix DFE-Platinum Series or N-SA devices with 128 MB of memory installed; or
  - 1 - 64 on Matrix DFE-Platinum Series or N-SA devices with an additional 128 MB memory upgrade.

Use the commands in the following sections to adjust these defaults.

Note: Poorly chosen adjustments to these parameters can have a negative impact on network performance. Please refer to the 1EEE 802.1D specification for guidance.

Setting the Bridge Priority

You can globally configure the priority of an individual bridge. When two bridges tie for position as the root bridge, this setting affects the likelihood that a bridge will be selected. The lower the bridge's priority, the more likely the bridge will be selected as the root bridge.

To set the bridge priority:

set spantree priority *priority* [*sid*]

Valid *priority* values are **0** - **61440** (in increments of 4096), with 0 indicating high priority and 61440 low priority.

Valid *sid* values are **0** - **4094**. If not specified, SID 0 will be assumed.

Setting a Port Priority

You can set a spanning tree priority for a port, which will be used to break the tie when two bridges tie for position as the root bridge. The bridge with the lowest port value will be elected.

To set a port priority:

set spantree portpri *port-string priority* [**sid** *sid*]

Valid *priority* values are **0** - **240** (in increments of 16) with 0 indicating high priority.

Valid *sid* values are **0** - **4094**. If not specified, SID 0 will be assumed.

Assigning Port Costs

Each interface has a spanning tree port cost associated with it, which helps to determine the quickest path between the root bridge and a specified destination. By convention, the higher the port speed, the lower the port cost. By default, this value is set to 0, which forces the port to recalculate spanning tree port cost based on the speed of the port and whether or not legacy (802.1D) path cost is enabled.

To assign different spanning tree port costs:

set spantree adminpathcost *port-string cost* [**sid** *sid*]

Va1id *cost* values are:

– **0** - **65535** if legacy path cost is enabled.

– **0** - **200000000** if legacy path cost is disabled.

Valid *sid* values are **0** - **4094**. If not specified, SID 0 will be assumed.

Notes: By default, legacy path cost is disabled. Enabling the device to calculate legacy path costs affects the range of valid values that can be administratively assigned.

To check the status of legacy path cost, use **show spantree legacypathcost**.

To disable legacy path cost, if necessary use **set spantree legacypathcost disable**.

Adjusting Bridge Protocol Data Unit (BPDU) Intervals

Use the commands in this section to adjust the following default BPDU interval values.

BPDU Interval	Default Value
Hello time (bridge and ports)	2 seconds
Forward delay	15 seconds
Maximum age time	20 seconds

Adjusting the Bridge Hello Time

Caution: Poorly chosen adjustments to bridge and port hello time parameters can have a negative impact on network performance. It is recommended that you do not change these parameters unless you are familiar with spanning tree configuration and have determined that adjustments are necessary. Please refer to the 1EEE 802.1D specification for guidance.

Hello time is the interval at which the bridge or individual ports send BPDU messages. By default, bridge hello mode is enabled, meaning the device uses a single bridge administrative hello time.

Adjust the bridge hello time as follows:

1. Check the status of bridge hello mode:

   `show spantree bridgehellomode`

2. If necessary, re-enable bridge hello mode:

   `set spantree bridehellomode enable`

3. Set a new hello time interval:

   `set spantree hello` *`interval`*

   Valid *interval* values are 1 - 10.

Adjusting Port Hello Times

You can set the device to use per-port administrative hello times by disabling bridge hello mode and adjusting the hello time interval for one or more ports as follows:

1. Check the status of bridge hello mode:

   `show spantree bridgehellomode`

2. If necessary, disable bridge hello mode:

   `set spantree bridehellomode disable`

3. Set a new hello time interval for one or more ports:

   `set spantree porthello` *`port-string interval`*

   Valid *interval* values are 10 - 100.

Adjusting the Forward Delay Interval

When rapid transitioning is not possible, forward delay is used to synchronize BPDU forwarding. The forward delay interval is the amount of time spent listening for topology change information after an interface has been activated for bridging and before forwarding actually begins. This delay is required because every device must receive information about topology changes before it starts to forward frames. In addition, each port needs time to listen for conflicting information that would make it return to a blocking state. Otherwise, temporary data loops might result.

To adjust the forward delay interval setting:

`set spantree fwddelay` *`delay`*

Valid *delay* values are 4 - 30.

Defining the Maximum Age

If a bridge does not hear BPDUs from the root bridge within a specified interval, it assumes that the network has changed and recomputes the spanning tree topology.

To adjust the maximum age setting:

set spantree maxage *agingtime*

Valid *agingtime* values are 6 - 40.

Setting the Maximum Configurable STPs

By default, the device is running in MSTP mode and, depending on the model type and memory installed, will allow you to configure up to 33 spanning tree instances.

To adjust the maximum configurable number of spanning trees allowed on the device:

set spantree maxconfigurablestps *numstps*

Valid *numstps* values are:

- 1-9 for Matrix DFE-Gold Series modules,
- 1-33 for Matrix DFE-Platinum Series or N-SA devices with128 MB of memory installed (device default), or
- 1-64 for Matrix DFE-Platinum Series or N-SA devices with and additional 128 MB memory upgrade installed.

Enabling the Backup Root Function

Disabled by default the backup root function works only when the backup root-enabled bridge is directly connected to the root bridge. It then prevents stale spanning tree information from circulating throughout the network in the event that the link between the root bridge and backup root-enabled bridge is lost. If this happens, the backup root will dynamically lower its bridge priority relative to the existing root bridge's priority, causing it to immediately be selected as the new root bridge.

To enable the backup root function on a SID:

set spantree backuproot *sid* **enable**

When SNMP trap messaging is configured and the backup root function is enabled, a trap message will be generated when the SID becomes the new root of the network.

Adjusting RSTP Parameters

Since rapid (one-second) link reconfiguration can happen only on a point-to-point link or an edge port (a port that is known to be on the edge of a bridged LAN), in some cases you may want to define them administratively. However, since edge port and point-to-point links are automatically detected on Matrix DFE Series and N-SA devices, in most cases you will not need to change these default port designations.

Defining Point-to-Point Links

By default, the administrative point-to-point status is set to auto on all spanning tree ports, allowing the Matrix DFE Series or N-SA firmware to determine each port's point-to-point status. In most cases, this setting will not need to be changed and will provide optimal RSTP functionality. You can, however, use the following commands to review and, if necessary, change the point-to-point status of a spanning tree link.

Review and define the point-to-point status of an RSTP link as follows:

1. Display the point-to-point operating status of a LAN segment attached to a port:

   show spantree operpoint [**port** *port-string*]

   A status of "true" indicates the LAN segment is operating as a point-to-point link.

   A status of "false" indicates it is not.

   If *port-string* is not specified, point-to-point operating status will be displayed for all spanning tree ports.

2. Display the point-to-point administrative status of a LAN segment attached to a port:

   show spantree adminpoint [**port** *port-string*]

   A status of "true" indicates the port is administratively set to be considered point-to-point.

   A status of "false" indicates the port is administratively set to be considered non point-to-point.

   A status of "auto" (the default setting) indicates that the firmware is allowed to determine the port's point-to-point status.

   If *port-string* is not specified, point-to-point administrative status will be displayed for all spanning tree ports.

3. If necessary, change the point-to-point administrative status of a LAN segment attached to a port:

   set spantree adminpoint *port-string* **auto** | **true**

Defining Edge Port Status

By default, edge port status is enabled on all ports, indicating that a port is on the edge of a bridged LAN. Automatic edge port detection is also enabled by default on the device. You can use the following commands to review and, if necessary, change the edge port detection status on the device and the edge port status of spanning tree ports.

Review and define edge port status as follows:

1. Display the status of edge port detection:

   show spantree autoedge

2. If necessary, re-enable edge port detection:

   set spantree autoedge enable

3. Display the edge port operating status of one or more port(s):

   show spantree operedge [**port** *port-string*]

   A status of "true" indicates the port is operating as an edge port.

   A status of "false" indicates it is not.

   If *port-string* is not specified, edge port status will be displayed for all spanning tree ports.

4. Display the edge port administrative status of one or more port(s):

   show spantree adminedge[**port** *port-string*]

   A status of "true" indicates the port is administratively set to be considered an edge port.

   A status of "false" indicates the port is administratively set to be considered a non edge port.

   If *port-string* is not specified, edge port administrative status will be displayed for all spanning tree ports.

5. If necessary, change the edge port administrative status of one or more port(s):

   set spantree adminedge *port-string* **true**

Configuring MSTP

In order for MSTP to provide multiple forwarding paths, the following must happen:

- The configuration identifier must match on all bridges within the region.
- All bridges must be within the same region.
- All bridges must be connected to MSTP-aware bridges. (They can be connected using a shared media such as a repeater provided that a single spanning tree device does not reside on that LAN).

Notes: A single spanning tree device between two MSTP bridges will terminate the ability to have multiple forwarding paths.

An MSTP bridge connected to a repeater will use timers to change port states and rapid transition will not occur. If the bridge must be connected to a repeater, **adminpoint** should be set to "true" on those ports connected to the repeater. For more information, refer back to Defining Point-to-Point Links on page 5-13.

This section provides information about the following MSTP tasks:

- Simple MSTP Configuration
- Adjusting MSTP Parameters
- Monitoring MSTP

Simple MSTP Configuration

The following example describes setting up the simple MSTP network shown in Figure 5-5. By default, each device will be in its own MST region using its own MAC address as the MST configuration ID. This configuration groups Device 1 and Device 2 into a single MST region with an MSTI configuration name of "South." It maps VLAN 2 to MSTI SID 2 and VLAN 3 to MSTI SID 3.

Figure 5-5 MSTP Sample Network Configuration

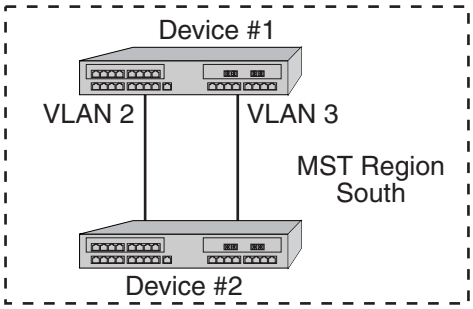

Procedure 5-1 shows how to configure Devices 1 and 2 for MSTP.

Procedure 5-1 Configuring Devices 1 and 2 for Simple MSTP

Step	Task	Command(s)
1.	Create VLANs 2 and 3.	**`set vlan create 2-3`**
2.	Set each device's configuration name to South.	**`set spantree mstcfgid South`**
3.	Create MSTI SID 2.	**`set spantree msti sid 2 create`**
4.	Create MSTI SID 3.	**`set spantree msti sid 3 create`**
5.	Create a FID-to-SID mapping for VLAN 2 to SID 2.	**`set spantree mstmap 2 sid 2`**
6.	Create a FID-to-SID mapping for VLAN 3 to SID 3.	**`set spantree mstmap 3 sid 3`**

Adjusting MSTP Parameters

You may need to adjust certain spanning tree parameters if the default values are not suitable for your bridge configuration. Refer back to Adjusting Spanning Tree Parameters (page 5-9) and Adjusting RSTP Parameters (page 5-12) for information on adjusting spanning tree defaults. Changes made to global and port-related spanning tree defaults will take affect if the device is running in STP, RSTP, or MSTP.

Monitoring MSTP

The commands listed in Table 5-4 allow you to monitor MSTP statistics and configurations on the Matrix DFE Series or N-SA device. You can also use the show commands described in "Reviewing and Re-enabling Spanning Tree" on page 5-8 to review information related to all spanning tree protocol activity.

Table 5-4 Commands for Monitoring MSTP

Task	Command
Verify that MSTP is running on the device.	**`show spantree version`**
Display the maximum configurable MSTIs allowed on the device.	**`show spantree maxconfigurablestps`**
Display a list of MSTIs configured on the device.	**`show spantree mstilist`**

Table 5-4 Commands for Monitoring MSTP (continued)

Task	Command
Display the mapping of one or more filtering database IDs (FIDs) to spanning trees. Since VLANs are mapped to FIDs, this shows to which SID a VLAN is mapped.	**show spantree mstmap** [**fid** *fid*]
Display the spanning tree ID(s) assigned to one or more VLANs.	**show spantree vlanlist** [*vlan-list*]
Display MST configuration identifier elements, including format selector, configuration name, revision level, and configuration digest.	**show spantree mstcfgid**
Display protocol-specific MSTP counter information.	**show spantree debug** [**port** *port-string*] [**sid** *sid*] [**active**]

Configuring Spanguard

This section provides information about the following Spanguard tasks:

- Overview of the Spanguard Function
- Enabling and Adjusting Spanguard

Overview of the Spanguard Function

Enterasys Networks' Spanguard function provides the ability for Enterasys switches to detect unauthorized bridges in the network. It protects against Spanning Tree Denial of Service (DoS) attacks as well as unintentional/unauthorized connected bridges. This is done by intercepting received BPDUs on configured ports and locking these ports so they do not process any received packets—thus protecting the integrity of the Spanning Tree topology.

By default, Spanguard is globally disabled. When enabled, reception of a BPDU on a port that is administratively configured as a spanning tree edge port (adminedge = True) will cause the port to become locked and the state set to blocking. When this condition is met, packets received on that port will not be processed for a specified timeout period. The port will become unlocked when either:

- The timeout expires
- The port is manually unlocked
- The port is no longer administratively configured as adminedge = True
- The Spanguard function is disabled

The port will become locked again should another offending BPDU be received on that port after expiration of the timeout or manual unlocking of that port occurs.

In the event of a DoS attack with Spanguard enabled and configured, no spanning tree topology changes or topology reconfigurations will be seen. The state of the spanning tree will be completely unaffected by the reception of any spoofed BPDUs regardless of the BPDU type, rate received or duration of the attack.

By default, when SNMP and Spanguard are enabled, a trap message will be generated when Spanguard detects that an unauthorized port has tried to join a Spanning Tree.

Enabling and Adjusting Spanguard

Use the following commands to configure device ports for Spanguard, to enable the Spanguard function, and to review Spanguard status on the device.

Reviewing and Setting Edge Port Status

Note: In order to utilize the Spanguard function, you must know which ports are connected between switches as ISLs (inter-switch links). Also, edge port status (adminedge = true or false) must be configured on the entire switch, as described in "Defining Edge Port Status" on page 5-13, before Spanguard will work properly.

Review and set edge port status as described in "Defining Edge Port Status" on page 5-13 as follows:

1. Review edge port administrative status on the device.
2. Set edge port administrative status to false on all known ISLs.
3. Set edge port administrative status to true on any remaining ports where Spanguard protection is desired. This indicates to Spanguard that these ports are not expecting to receive any BPDUs. If these ports do receive BPDUs, they will become locked.

Enabling and Adjusting Spanguard

Use the commands in this section to enable Spanguard, to review Spanguard status, as necessary, and to adjust Spanguard parameters.

To enable Spanguard on the device:

```
set spantree spanguard enable
```

Use this command to adjust the Spanguard timeout value. This sets the length of time that a Spanguard-affected port will remain locked:

```
set spantree spanguardtimeout timeout
```

Valid values are 0 - 65535 seconds. Default is 300 seconds. Setting the value to 0 will set the timeout to forever.

Use this command to manually unlock a port that was locked by the Spanguard function. This overrides the specified timeout variable:

```
set spantree spanguardlock port-string
```

Reviewing Spanguard Status and Settings

Use the following commands to review Spanguard status and settings.

To verify Spanguard status:

```
show spantree spanguard
```

To review the status of the Spanguard lock function on one or more ports:

```
show spantree spanguardlock [port port-string]
```

If not specified, status for all ports will be displayed.

To review the Spanguard timeout setting on one or more ports:

```
show spantree spanguardtimeout
```

To review the status of the Spanguard trap function:

```
show spantree spanguardtrapenable
```

6

Managing Syslog

This chapter provides information about the following system logging procedures on the Matrix DFE Series and N-SA devices.

For information about...	Refer to page...
Logging Overview	6-1
Interpreting Messages	6-2
Configuring Syslog Servers, Applications, and Console Logging	6-3

Logging Overview

The Syslog implementation on the Matrix DFE Series or N-SA device uses a series of system logging messages to track device activity and status. These messages inform users about simple changes in operational status or warn of more severe issues that may affect system operations. Logging can be configured to display messages at a variety of different severity levels about application-related error conditions occurring on the device.

As defined in RFC 3164, the Syslog protocol allows administrators to place messages into one of several broad categories. These broad categories generally consist of the facility that generated them, along with an indication of the severity of the message. This helps administrators prioritize and selectively filter messages, while also giving them the ability to place status or informative messages in a file for later review. Administrators can decide to have all messages stored locally as well as to have all messages of a high severity forwarded to another device. They can also have messages from a particular facility sent to some or all of the users of the device, and displayed on the system console. However you decide to configure the disposition of the event messages, the process of having them sent to a Syslog collector generally consists of:

- Determining which facility messages and which severity levels will be forwarded.
- Defining the remote receiver, also known as the Syslog server.

For example, you may want all messages that are generated by the mail facility to be forwarded to one particular Syslog server.

Use the commands in this chapter to complete these tasks and customize your logging configuration.

Syslog Terms and Definitions

Table 6-1 lists terms and definitions used in Syslog configuration.

Table 6-1 Syslog Terms and Definitions

Term	Definition
Facility	The Syslog specification uses a facility code to categorize which functional process is generating an error message. Syslog combines the facility and severity values to determine message priority. The Matrix DFE Series and N-SA implementation uses the eight facility designations reserved for local use: **local0** - **local7**. Default is **local4**. For more information about facility designations, refer to RFC 3164.
Severity	A value used to indicate the severity of the error condition generating the Syslog message. The Matrix DFE Series and N-SA implementation provides the following 8 levels: **1** - emergencies (system is unusable) **2** - alerts (immediate action required) **3** - critical conditions **4** - error conditions **5** - warning conditions **6** - notifications (significant conditions) **7** - informational messages **8** - debugging messages
Application	Client process for which Syslog is tracking error conditions. Supported applications and their associated CLI mnemonic values are: **RtrAcl** - Access Control Lists **CLI** - Command Line Interface **SNMP** - Simple Network Management Protocol **Webview** - Enterasys Web-based system management **System** - System messages **RtrFe** - Router Forwarding Engine **Trace** - Trace logging **RtrLSNat** - Load Share Network Address Translation **FlowLimt** - Flow limiting **UPN** - User Personalized Networks
Syslog server	A remote server configured to collect and store Syslog messages. The Matrix DFE Series and N-SA implementation allows up to 8 server IP addresses to be configured for Syslog.

Interpreting Messages

Every system message generated by the Matrix DFE Series or N-SA device follows the same basic format:

```
<facility/severity> <time stamp> <address> <application> <message text>
```

Example

This example shows Syslog informational messages, displayed with the **show logging buffer** command. It indicates that messages were generated by facility 16 at severity level 5 from the CLI application on IP address 10.42.71.13. Table 6-2 describes the components of these messages.

```
Matrix(rw)->show logging buffer

<165>Sep  4 07:43:09 10.42.71.13 CLI[5]User:rw logged in from 10.2.1.122
(telnet)
<165>Sep  4 07:43:24 10.42.71.13 CLI[5]User: debug failed login from
10.4.1.100
(telnet)
```

Table 6-2 Syslog Message Components

Component	Description	Example Code
Facility/Severity	Combined code indicating the facility generating the message and the severity level. Facility codes 16 - 23 are Syslog designations for local0 - local7. For a complete list of facility codes, refer to RFC 3164.	`<165>` = A message from facility local0 (numerical code 16) at severity 5.
Time stamp	Month, date, and time the Syslog message appeared.	`Sep  4 07:43:09`
Address	IP address of the client originating the Syslog message.	`10.42.71.13`
Application	Client process generating the Syslog message.	`CLI`
Slot#	Slot location of the device module generating the Syslog message.	`(5)` = Slot 5 in the chassis.
Message text	Brief description of error condition.	`User: debug failed login from 10.4.1.100 (telnet)`

Configuring Syslog Servers, Applications, and Console Logging

The Matrix DFE Series or N-SA device uses console and Syslog server logging to allow you to monitor and track system activity. Messages reporting various states and events can be viewed and discarded or stored for future reference. By default, Syslog is configured to log messages for all applications listed in Table 6-1 at a severity level of 6 (warnings of significant conditions.) and a local4 facility category. Syslog server is globally enabled, with no IP addresses configured, at a severity level of 8.

Use the procedures in this section to perform the following logging configuration tasks:

- Enabling Syslog Server(s)
- Modifying Syslog Server Defaults
- Reviewing and Configuring Logging for Applications
- Enabling Console Logging and File Storage

About Server Versus Application Severity Levels

The default Syslog configuration allows client applications to generate messages on a severity level of 6 and the Syslog server to log messages from clients at a severity level of 8. This means that all enabled servers will accept messages from all logging applications generated for error conditions at levels 6, 7 and 8. You can use the procedures described in this chapter to change these parameters, modifying the logging behavior between one or more client applications and one or more servers.

Enabling Syslog Server(s)

Use the following commands to enable Syslog on one or more servers and verify the configuration:

1. Enable Syslog on a server IP:

   set logging server *index* **ip-addr** *ip-addr* **state enable**

   Index is a value from 1 to 8 that specifies the server table index number for this server.

2. (Optional) Verify the server configuration:

   show logging server [*index*]

   If *index* is not specified, information for all configured Syslog servers will be displayed.

Example

This sample output from the **show logging server** command shows that Syslog has been enabled on two servers. These servers are using the default UDP port 514 to receive messages from clients and are configured to log messages from the local1 and local2 facilities, respectively. Logging severity on both servers is set at 5, which will record warning conditions. These settings can be changed on a per-server basis, or for all servers using the commands described in the next section.

```
Matrix(rw)->show logging server

   IP Address     Facility   Severity     Description    Port     Status
-------------------------------------------------------------------------
 1 132.140.82.111 local1     warning(5)    default         514    enabled
 2 132.140.90.84  local2     warning(5)    default         514    enabled
```

Modifying Syslog Server Defaults

Unless otherwise specified, servers enabled for Syslog will automatically adopt factory default settings for facility name and description, message severity level, and UDP port number. Use the following commands to change these settings either during or after enabling a new server.

Displaying system logging defaults

To display system logging defaults, or all logging information, including defaults:

show logging {**default**|**all**}

Modifying default settings

You can change factory default logging settings using one of the following methods.

- To specify logging parameters during or after new server setup:

  set logging server *index* **ip-addr** *ip-addr* [**facility** *facility*] [**severity** *severity*] [**descr** *descr*] [**port** *port*] **state enable**

  If not specified, optional server parameters will be set to the following system defaults:

  – *facility* - local4

  – *severity* - 8 (accepting all messages above and including the level configured for applications).

  – *descr* - no description applied

  – *port* - UDP port 514

- To change default parameters for all servers:

  set logging default {[**facility** *facility*] [**severity** *severity*] [**port** *port*]}

Examples

This example shows how to configure Syslog server 1 at IP 134.141.89.113 to accept messages from facility category local6 at severity level 3 and lower.

```
Matrix(rw)->set logging server 1 ip-addr 134.141.89.113 facility local6
severity 3
```

This example shows how to change Syslog defaults so that all servers accept messages from the local2 facility category at a severity level 4 and lower. These settings will apply to all newly-enabled servers, unless explicitly configured with the **set logging server** command:

```
Matrix(rw)->set logging default facility local2 severity 4
```

Reviewing and Configuring Logging for Applications

By default, messages will be logged at severity level 6 to configured servers, and to the console and persistent file (if enabled) for all applications running on your device. Use the commands in this section to review and change logging settings for all current applications.

Note: The logging severity for client applications must be set at the same level — or a lower level — than the severity level set for the Syslog server in order for the server to collect error messages from those clients. For example, a Syslog server with a severity level 4 would not collect application-related messages from a client set at level 5 or above.

Displaying Current Application Severity Levels

To display logging severity levels for one or all applications currently running on your device:

show logging application {*mnemonic*|**all**}

Example

This example shows output from the **show logging application all** command. A mnemonic value for each application is listed with the severity level at which logging has been

configured and the server(s) to which messages will be sent. In this case, logging for applications has not been changed from the default severity level of 6:

```
Matrix(rw)->show logging application all
 Global Logging State: Enabled

        Application   Current Severity Level Server
-------------------------------------------------------
 88     RtrAcl                    6              1-8
 89     CLI                       8              1-8
 90     SNMP                      8              1-8
 91     Webview                   6              1-8
 93     System                    6              1-8
 95     RtrFe                     6              1-8
 96     Trace                     6              1-8
 105    RtrLSNat                  6              1-8
 111    FlowLimt                  6              1-8
 112    UPN                       6              1-8

1(emergencies)  2(alerts)       3(critical)
4(errors)       5(warnings)     6(notifications)
7(information)  8(debugging)
```

Note: Mnemonic values are case sensitive and must be typed as they are listed in Table 6-1.

Modifying Severity Levels and Assigning Syslog Servers for Applications

By default, messages will be generated by client applications at a severity level of 6 and logged to enabled Syslog servers at the system's default severity level of 8, or to a level set using the **set logging server** or **set logging default** command. To modify the severity level of log messages and the server(s) to which messages will be sent for one or all applications:

set logging application {[*mnemonic*|**all**]}[**level** *level*][**servers** *servers*]

Example

This example shows how to set the severity level for SSH (Secure Shell) to 5 so that warning conditions will be generated for that application and sent to Syslog server 1:

```
Matrix(rw)->set logging application SSH level 5 server 1
```

Note: The severity value of the specified server(s) must be configured the same — or at a higher value — as the severity specified for applications in order for those applications to be logged to the server(s).

Enabling Console Logging and File Storage

You can configure logging to display messages to the current console CLI session only, or to display to the console and save to a persistent file. Console logging allows you to view only as many messages as will fit on the screen. As new messages appear, old messages simply scroll off the console. While this is a temporary means of logging information, it allows administrators to track very specific activities quickly and easily. Console log messages can also be saved to a persistent file at two locations:

- slotX/logs/current.log — Location of current system log messages (up to 256k), where X specifies the slot location of the device.
- slotX/logs/old.log — Location of previous system log messages, where X specifies the slot location of the device. Current messages will be moved to the old.log when current.log file exceeds 256k.

Use the following commands to review and configure console logging and file storage.

Displaying to the console and saving to a file

To display log messages to the console and save to a persistent file:

```
set logging local console enable file enable
```

Displaying to the current CLI session

To display logging to the current CLI console session:

```
set logging here enable
```

This adds the current CLI session to the list of Syslog destinations, and will be temporary if the current CLI session is using Telnet or SSH.

Displaying a log file

To display the contents of the persistent log file:

show file slot*slotnumber*/**logs**/**current.log**|**old.log**

Note: These log files may also be copied to another device using FTP or TFTP.

7

Configuring IP

This chapter provides information about the following Internet Protocol (IP) configuration procedures on the Matrix DFE Series and N-SA devices. These procedures involve general non-routing protocol configuration tasks that are independent of routing protocol (such as OSPF) operation of the device.

For information on configuring the OSPF routing protocol, refer to Chapter 8, Configuring OSPF.

For information on configuring the VRRP routing protocol, refer to Chapter 9, ***Configuring VRRP***.

For information about...	Refer to page...
Using the Matrix DFE Series or N-SA Device as a Router	7-1
Reviewing and Configuring Router Interfaces	7-5
Managing Router Configuration Files	7-10
Reviewing and Configuring the ARP Table	7-12
Configuring IP Broadcast Settings	7-14
Configuring Routes and Monitoring IP Traffic	7-16
Configuring ICMP	7-18

Using the Matrix DFE Series or N-SA Device as a Router

While Matrix DFE Series and N-SA devices operating in switch mode allow users within a configured VLAN to communicate, routing capabilities are necessary to transmit traffic from VLAN to VLAN. This is the purpose of configuring a Matrix DFE Series or N-SA device for router operation: to allow users on different VLANs to communicate with each other. Reviewing and Configuring Router Interfaces (page 7-5) describes how to configure a switch for routing to serve that purpose, and provides a sample configuration. This section also provides the following information about using the device as a router:

- Example Scenario
- Preparing the Device for Router Mode
- Pre-Routing Configuration Tasks
- Enabling Router Configuration Modes

Example Scenario

Imagine you are using a Matrix DFE Series or N-SA switch device that is configured for 10 VLANs. Each VLAN includes 20 physical ports with one client on each port, for a total of 200 users in the 10 VLANs. Since, in the switching world, the only users that can talk to each other are the users in the same VLAN, the 20 users in VLAN 3 can talk to each other, but they can't talk to the 180 other users in the rest of your network.

This is where routing gets involved. For traffic to pass from a host in one VLAN to a host in another VLAN, the frames must be processed by a router, and then forwarded directly to the destination host (if that host happens to be directly connected to the router), or forwarded to the next router in the network on the way to the target host.

Before traffic can be processed in this way, a relationship must exist between the router being configured and the underlying switch. This relationship is established with the VLAN configuration described in the Sample Configuration on page 7-8. Every switched VLAN in which users are to be granted access to other VLANs, must have a corresponding router "interface VLAN" configured. The procedures necessary to complete this process are detailed in the following sections.

Preparing the Device for Router Mode

At least one module must be configured to run as a router to operate a Matrix DFE Series or N-SA device in router mode. Once the tasks described in "Pre-Routing Configuration Tasks" (page 7-3) are completed, you can follow the procedures in this chapter to perform basic IP interface configuration tasks, and to set the general, non-protocol-specific routing parameters required for IP routing.

About Loopback and VLAN Interfaces

The Matrix DFE Series device allows you to configure both VLAN and loopback interfaces for IP routing. Loopback interfaces are different from VLAN routing interfaces because they allow you to disconnect the operation of routing protocols from network hardware operation, improving the reliability of IP connections. A loopback interface is always reachable. The IP address assigned to the loopback interface is used as the router ID, which helps when running protocols like OSPF, because OSPF can be running even when the outbound interface is down. IP packets routed to the loopback interface are rerouted back to the router or access server and processed locally.

For an example of how a loopback interface could be used in a routing configuration, refer to Chapter 8, Configuring OSPF.

Routing interface configuration commands in this guide will configure either a VLAN or loopback interface, depending on your choice of parameters, as shown in Table 7-1.

Table 7-1 VLAN and Loopback Interface Configuration Modes

For Routing Interface Type...	Enter (in Global Configuration Mode)...	Resulting Prompt...
VLAN	**vlan** *vlan-id*	Matrix>Router1 (config-if(Vlan 1))#
Loopback	**loopback** *loopback-id*	Matrix>Router1 (config-if (Lpbk 1))#

For details on how to enable all router CLI configuration modes, refer to Table 7-2 on page 7-3.

For information on configuring the OSPF routing protocol, refer to Chapter 8.

For information on configuring the VRRP routing protocol, refer to Chapter 9.

Pre-Routing Configuration Tasks

The discussed previously, following pre-routing tasks must be performed from the switch CLI.

- Starting the Command Line Interface (page 2-1)
- Setting System Passwords (page 2-4)
- Setting and Verifying a System IP Address (page 2-4)
- Setting the Time, Date, and System Information (page 2-4)
- Creating and enabling VLANs
- Managing Configuration and Image Files (page 2-7)
- Downloading a New Firmware Image (page 2-8)
- Activating Licensed Features (page 2-13)
- Preparing the Device for Router Mode (page 3-16)

Enabling Router Configuration Modes

The Matrix DFE Series / N-SA CLI provides different modes of router operation for issuing a subset of commands from each mode. Table 7-2 describes these modes of operation. Modes in the shaded area will be referred to in this chapter. Modes in non-shaded areas will be referred to in later chapters covering routing protocols, such as RIP and OSPF.

Notes: Sample prompts shown in Table 7-2 and throughout this chapter assume module 1 (or the N-SA device) has been configured for routing, and that VLAN or loopback 1 will be designated as routing interfaces.

A *module* designation of 1 must be entered to enable routing on the Matrix N-SA standalone device. All other values will result in an error message.

Table 7-2 Router CLI Configuration Modes

Use this mode...	To...	Access method...	Resulting Prompt...
Privileged EXEC Mode	Set system operating parameters Show configuration parameters Save/copy configurations	From the switch CLI: 1. Type **router** *module* (using a module number configured for routing). 2. Type **enable**.	`Matrix>Router1>` `Matrix>Router1#`
Global Configuration Mode	Set system-wide parameters.	Type **configure terminal** from Privileged EXEC mode.	`Matrix>Router1(config)#`
Interface Configuration Mode	Configure router interfaces.	Type **interface vlan** or **interface loopback** and the interface's *id* from Global Configuration mode.	`Matrix>Router1 (config-if(Vlan 1 | Lpbk 1))#`

Table 7-2 Router CLI Configuration Modes (continued)

Use this mode...	To...	Access method...	Resulting Prompt...
Router Configuration Mode	Set IP protocol parameters.	Type **router** and the *protocol name* from Global or Interface Configuration mode.	`Matrix>Router1 (config-router)#`
Key Chain Configuration Mode	Set protocol (RIP) authentication key parameters.	Type **key chain** and the key chain *name* from Router (RIP) Configuration mode.	`Matrix>Router1 (config-keychain)#`
Key Chain Key Configuration Mode	Configure a specific key within a RIP authentication key chain.	Type **key** and the *key-id* from Key Chain Configuration Mode.	`Matrix>Router1 (config-keychain-key)#`
Policy-Based Routing Configuration Mode	Configure policy-based routing for a route map.	Type **route-map**, an *id-number*, and **permit** or **deny** from Global Configuration Mode.	`Matrix>Router1 (config-route-map-pbr)#`
Server Load Balancing (SLB) Server Farm Configuration Mode	Configure an LSNAT server farm.	Type **ip slb serverfarm** and the *serverfarmname* from Global Configuration Mode.	`Matrix>Router1 (config-slb-sfarm)#`
Server Load Balancing (SLB) Real Server Configuration Mode	Configure an LSNAT real server.	Type **real** and the real server *IP address* from SLB Server Farm Configuration Mode.	`Matrix>Router1 (config-slb-real)#`
Server Load Balancing (SLB) Virtual Server Configuration Mode	Configure an LSNAT virtual server.	Type **ip slb vserver** and the *vserver-name* from Global Configuration Mode.	`Matrix>Router1 (config-slb-vserver)#`
IP Local Pool Configuration Mode	Configure a local address pool as a DHCP subnet	Type **ip local pool** and the local pool *name* from Global Configuration Mode.	`Matrix>Router1 (ip-local-pool)#`
DHCP Pool Configuration Mode	Configure a DHCP server address pool.	Type **ip dhcp pool** and the address pool *name* from Global Configuration Mode.	`Matrix>Router1 (config-dhcp-pool)#`
DHCP Class Configuration Mode	Configure a DHCP client class.	Type **client-class** and the client class *name* from DHCP Pool or Host Configuration Mode.	`Matrix>Router1 (config-dhcp-class)#`

Table 7-2 Router CLI Configuration Modes (continued)

Use this mode...	To...	Access method...	Resulting Prompt...
DHCP Host Configuration Mode	Configure DHCP host parameters.	Type **client-identifier** and the *identifier*, or **hardware-address** and an *address* from any DHCP configuration mode.	`Matrix>Router1 (config-dhcp-host)#`

Note: To jump to a lower configuration mode, type **exit** at the command prompt. To revert back to switch CLI, type **exit** from Privileged EXEC router mode.

Reviewing and Configuring Router Interfaces

Use the procedures in this section to perform the following router interface configuration tasks:

- Displaying Interface Statistics and Settings
- Configuring Interfaces for IP Routing

Displaying Interface Statistics and Settings

From any router mode, use this command to display information about one or more interfaces (VLANs or loopbacks) configured on the router:

show [**ip**] **interface** [**vlan** *vlan-id* | **loopback** *loopback-id*]

Specifying **ip** will display IP-related information, including status and configuration details, about routing interfaces. If not specified, statistics for all interfaces configured for routing will be displayed.

If interface type and ID are not specified, information for all known interfaces will be displayed.

Examples

This example shows how to use the **show interface** command to display information for all interfaces configured on the router. In this case, one loopback interface has been configured for routing. Table 7-3 provides a detailed description of the command output:

```
Matrix>Router1#show interface
Loopback 1 is   Administratively UP
 Loopback 1 is Operationally UP
 Internet Address is 21.1.1.1, Subnet Mast is 255.255.255.0
 The name of this device is Loopback 1
 The MTU is 1500 bytes
 The bandwidth is 10000 Mb/s
```

This example shows how to use the **show ip interface** command to display IP configuration information for VLAN 1, including administrative status, IP address, MTU (Maximum Transmission Unit) size and bandwidth, and ACL configurations. Table 7-3 provides a detailed description of the command output:

```
Matrix>Router1#show ip interface vlan 1
 Vlan 1 is Admin UP
 Vlan 1 is Oper UP
 IP Address 81.81.7.3 Mask 255.255.255.128
 Frame Type ARPA
 MAC-Address 0001.f4da.16ee
 Incoming Access List is not Set
 Outgoing Access List is not Set
 IP Helper Address is not Set
 MTU is 1500 bytes
 ARP Timeout is 14400 seconds
 Proxy Arp is Enabled
 ICMP Re-Directs are enabled
 ICMP Unreachables are always sent
 ICMP Mask Replies are always sent
 Policy routing enabled, route map 101
```

Table 7-3 show ip interface Command Output

Output	What It Displays...
Vlan | Lpbk *N*	Whether the interface is administratively and operationally up or down.
IP Address	Interface's IP address and mask. Set using the **ip address** command as described in "Assigning an IP Address" on page 7-7.
Frame Type	Encapsulation type used by this interface. ARPA is the device default, and cannot be changed.
MAC-Address	MAC address mapped to this interface. Set using the **ip mac-address** command as described in "Assigning a MAC Address to an Interface" on page 7-14.
Incoming | Outgoing Access List	Whether or not an access control list (ACL) has been configured on this interface.
IP Helper Address	Whether or not an IP address has been designated for forwarding UDP datagrams from this interface. Set using the **ip helper-address** command as described in "Enabling or Disabling UDP Broadcast Forwarding" on page 7-15.
MTU	Interface's Maximum Transmission Unit size.
ARP Timeout	Duration for entries to stay in the ARP table before expiring. Default of 14,400 seconds can be modified using the **arp timeout** command as described in "Changing the ARP Timeout" on page 7-14.
Proxy Arp	Whether or not proxy ARP is enabled or disabled for this interface. Default state of enabled can be modified using the **ip proxy arp** command as described in "Disabling or Re-enabling Proxy ARP on an Interface" on page 7-14.
ICMP	ICMP (ping) settings. By default, ICMP messaging is enabled on a routing interface for both echo-reply and mask-reply modes. If, for security reasons, ICMP has been disabled, it can be re-enabled using the **ip icmp** command as described in "Enabling or Disabling ICMP" on page 7-18.
Policy routing	Whether or not policy-based routing has been configured on this interface.

Configuring Interfaces for IP Routing

Each Matrix DFE Series or N-SA module can support up to 256 routing interfaces, which can be configured for IP routing using the procedures in this section.

Creating and Enabling IP Routing Interfaces

Each VLAN or loopback interface must be configured for routing separately using the **interface** command. This command is also used to enable interface configuration mode, which is required for completing interface-specific configuration tasks.

To end configuration on one interface before configuring another, type **exit** at the command prompt.

From global configuration mode, use this command to create and/or enable an interface for IP routing. This command will enable interface configuration mode from global configuration mode, and, if the interface has not previously been created, this will create a new routing interface.

interface {vlan *vlan-id* | **loopback** *loopback-id*}

Automatically enabling an interface for routing at device startup also requires configuration with the **no shutdown** command as described below.

Assigning an IP Address

From interface configuration mode, use this command to assign an IP address to a routing interface.

ip address *ip-address ip-mask* [**secondary**]

If specified, the address will be assigned as a **secondary** IP address to the interface. Otherwise, it will be the interface's primary IP address. Each Matrix DFE Series or N-SA routing module supports up to 50 secondary addresses for each primary IP address, with up to a maximum of 200 secondary addresses allowed per router.

Auto-Enabling Routing Interfaces

From interface configuration mode, use this command to allow an interface to automatically be enabled at device startup:

no shutdown

Removing an IP Address and Disabling Routing

From interface configuration mode, use this "no" form of the **ip address** command to remove the specified IP address from the interface, and disable the interface for IP processing:

no ip address *ip-address ip-mask*

Alternatively, you can also use the **shutdown** command from interface configuration mode to disable IP routing on the interface.

Sample Configuration

The following sample configuration assumes the example network scenario described in "Using the Matrix DFE Series or N-SA Device as a Router" on page 7-1. The administrator wants to allow five out of 10 VLANs to exchange traffic. These are VLANs 20, 30, 40, 50, and 60. Procedure 7-1 lists the steps and the associated commands necessary to complete this process.

To see how this configuration would look in the CLI, refer to the example on page 7-8.

For a description of all the router configuration modes used in the Matrix DFE Series / N-SA CLI, refer back to Table 7-2.

Procedure 7-1 Configuring VLANs for Routing

Step	Task	Command(s)
Complete these steps for each of the network's 10 VLANs		
1.	**In switch mode:** Create VLANs.	**set vlan create** *vlan-list*
2.	Set the inbound port/VLAN (PVID) associations for untagged frames for each VLAN.	**set port vlan port-string pvid**
3.	Add PVID ports to the egress list of each VLAN, and allow them to transmit untagged frames.	**set vlan egress** *vlan-list port-string* **untagged**
Complete these steps to configure the Matrix DFE or N-SA module for routing		
4.	Configure a routing module.	**set router** *module*
5.	Enable router mode.	**router** *module*
6.	**In router mode:** Enable router Privileged EXEC mode.	**enable**
7.	**In router Privileged EXEC mode**: Enable global router configuration mode.	**configure terminal**
Complete these steps to assign routing interfaces to each of the five VLANs being configured for routing		
8.	**In Global Configuration mode:** Enable interface configuration mode using the interface of the routing module.	**interface vlan** *vlan-id*
9.	**In Interface Configuration mode:** Assign an IP address to the routing interface.	**ip address** {*ip-address ip-mask*}
10.	Enable the interface for IP routing.	**no shutdown**

Example

This example shows the complete configuration described in Procedure 7-1.

```
**Create the VLANs**
Matrix(rw)->set vlan create 10,20,30,40,50,60,70,80,90,100
```

```
**Set the inbound PVID associations for untagged frames**
Matrix(rw)->set port vlan fe.1.1-20 10
Matrix(rw)->set port vlan fe.1.21-40 20
Matrix(rw)->set port vlan fe.2.1-20 30
Matrix(rw)->set port vlan fe.2.21-40 40
Matrix(rw)->set port vlan fe.3.1-20 50
Matrix(rw)->set port vlan fe.3.21-40 60
Matrix(rw)->set port vlan fe.4.1-20 70
Matrix(rw)->set port vlan fe.4.21-40 80
Matrix(rw)->set port vlan fe.5.1-20 90
Matrix(rw)->set port vlan fe.5.21-40 100

**Set the list of egress ports and tagging for the VLANs**
Matrix(rw)->set vlan egress 10 fe.1.1-20 untagged
Matrix(rw)->set vlan egress 20 fe.1.21-20 untagged
Matrix(rw)->set vlan egress 30 fe.2.1-20 untagged
Matrix(rw)->set vlan egress 40 fe.2.21-20 untagged
Matrix(rw)->set vlan egress 50 fe.3.1-20 untagged
Matrix(rw)->set vlan egress 60 fe.3.21-20 untagged
Matrix(rw)->set vlan egress 70 fe.4.1-20 untagged
Matrix(rw)->set vlan egress 80 fe.4.21-20 untagged
Matrix(rw)->set vlan egress 90 fe.5.1-20 untagged
Matrix(rw)->set vlan egress 100 fe.5.21-20 untagged

**Configure the router**
Matrix(rw)->set router 1
Matrix(rw)->router 1
Matrix>Router1>enable
Matrix>Router1#configure terminal

**Assign and enable the routing interfaces**
Matrix>Router1(config)#interface vlan 20
Matrix>Router1(config-if(Vlan 20))#ip address 100.20.20.1 255.255.255.0
Matrix>Router1(config-if(Vlan 20))#no shutdown
Matrix>Router1(config-if(Vlan 20))#exit
Matrix>Router1(config)#interface vlan 30
Matrix>Router1(config-if(Vlan 30))#ip address 100.30.30.1 255.255.255.0
Matrix>Router1(config-if(Vlan 30))#no shutdown
Matrix>Router1(config-if(Vlan 30))#exit
Matrix>Router1(config)#interface vlan 40
Matrix>Router1(config-if(Vlan 40))#ip address 100.40.40.1 255.255.255.0
Matrix>Router1(config-if(Vlan 40))#no shutdown
Matrix>Router1(config-if(Vlan 40))#exit
Matrix>Router1(config)#interface vlan 50
Matrix>Router1(config-if(Vlan 50))#ip address 100.50.50.1 255.255.255.0
Matrix>Router1(config-if(Vlan 50))#no shutdown
Matrix>Router1(config-if(Vlan 50))#exit
Matrix>Router1(config)#interface vlan 60
Matrix>Router1(config-if(Vlan 60))#ip address 100.20.20.1 255.255.255.0
Matrix>Router1(config-if(Vlan 60))#no shutdown
```

Since IP routing is enabled by default when the router is “set”, the users on these five locally connected VLANs can now talk to each other, but not to the other five in the netowork. If the users in these five networks want to talk to users in other networks on other routers, then a routing protocol must be configured, such as OSPF.

For information on configuring the OSPF routing protocol, refer to Chapter 8.

Managing Router Configuration Files

Each Matrix DFE Series or N-SA device provides a single configuration interface which allows you to perform both switch and router configuration with the same command set. This section demonstrates managing configuration files while operating in router mode only. For a sample of how to use these commands interchangeably with the Matrix DFE Series / N-SA single configuration interface commands, refer to "Managing Configuration and Image Files" (page 2-7).

Use the procedures in this section to perform the following IP configuration file management tasks:

- Displaying the Running Configuration
- Saving or Erasing the Running Configuration
- Removing the Routing Configuration from the Device
- Performing a Basic Router Configuration

Displaying the Running Configuration

From any router mode, use this command to display the non-default, user-supplied commands entered while configuring the device:

```
show running-config
```

Example

This example shows how to display the current router operating configurativen:

```
Matrix>Router1#show running-config
!
license advanced abcdefg123456789
!
 Router id 182.127.62.1
!
route-map 102 permit 1
  match ip address 2 4 6 8 110
  set next-hop 10.2.1.1 10.2.2.1 10.2.3.1
!
interface vlan 2
  ip address  2.1.1.1  255.255.255.0
  ip policy route-map 102
  ip policy priority only
  ip policy loadpolicy round-robin
  ip policy pinger on
  ip policy pinger interval 10
  ip policy pinger retries 4
        no shutdown
!
 router rip
  network 182.127.0.0
```

Saving or Erasing the Running Configuration

From privileged EXEC mode, use this command to save the router running configuration to NVRAM, to erase it, or to display it to output devices.

write [**erase** | **file** [**filename** *config-file*] | **terminal**]

If **filename** *config-file* is not specified, the configuration will be saved to startup.cfg.

If no options are specified, the configuration will be displayed to the terminal.

Note: The **write file** command must be executed in order to save the router configuration to NVRAM. If this command is not executed, router configuration changes will not be saved upon reboot.

Removing the Routing Configuration from the Device

From global configuration mode, use this command to disable IP routing on the device and remove the routing configuration. By default, IP routing is enabled when interfaces are configured for it as described in Creating and Enabling IP Routing Interfaces on page 7-7.

no ip routing

Performing a Basic Router Configuration

Using Router-Only Config Files

Although the Matrix DFE Series / N-SA's single configuration interface provides one set of commands to perform both switch and router configuration, it is still possible to use router-only commands to configure the router. To do so, you need to add router config wrappers to your existing router config files, as shown in the following example.

Example

This example shows how to add router config wrappers to a file. In this case, "begin router 1" is placed at the beginning, and "end router 1" was placed at the end of the config file:

```
begin router 1

enable
conf t
write file
exit
disable
exit

end router 1
```

Displaying or Writing the Current Config to a File

The Matrix DFE Series / N-SA's single configuration interface allows you use the **show config** command to display or write the current router configuration to a file. For details, refer back to Displaying the Running Configuration on page 7-10.

Configuring the Router

You can configure the router using either of the following methods.

Using a downloaded file...

1. Download a router config file to the chassis using the **copy** command
2. Run the **configure** command using the downloaded config file.

For more information on these commands, refer to "Managing Configuration and Image Files" (page 2-7).

Creating and saving a custom file...

1. Configure a module for routing using the **set router** command as described in Procedure 3-1.
2. Enable the router as described in Procedure 3-1 and configure it manually.
3. Save the configuration using the **write file** command as described in "Saving or Erasing the Running Configuration" (page 7-11).

Reviewing and Configuring the ARP Table

The Matrix DFE Series or N-SA device allows you to configure Address Resolution Protocol (ARP) table entries and parameters. ARP is used to associate IP addresses with MAC addresses. Once determined, the IP address/MAC association is stored in an ARP cache for rapid retrieval. An IP datagram is then encapsulated into a link-layer frame and sent over the network.

Use the procedures in this section to perform the following ARP table configuration tasks:

- Displaying the ARP Table
- Adding or Removing Static ARP Entries
- Disabling or Re-enabling Proxy ARP on an Interface
- Assigning a MAC Address to an Interface
- Changing the ARP Timeout
- Clearing the ARP Cache

Displaying the ARP Table

From any router mode, use this command to display entries in the ARP table.

show ip arp [*ip-address*] [**vlan** *vlan-id*] [*output-modifier*]

Optional output modifiers are as follows:

- | **begin** *ip-address* — Displays only ARP entries that begin with the specified IP address.
- | **exclude** *ip-address* — Excludes ARP entries matching the specified IP address.
- | **include** *ip-address* — Includes ARP entries matching the specified IP address.

If no parameters are specified, all entries in the ARP table will be displayed.

Examples

This example shows how to display all entries in the ARP table:

```
Matrix>Router1#show ip arp
Protocol   Address          Age (min) Hardware Addr   Type  Interface
-------------------------------------------------------------------------
Internet  134.141.235.251   0         0003.4712.7a99  ARPA   Vlan1
Internet  134.141.235.165   -         0002.1664.a5b3  ARPA   Vlan1/fe.1.1
Internet  134.141.235.167   4         00d0.cf00.4b74  ARPA   Vlan
```

This example shows how to display an entry related to a specific IP address:

```
Matrix>Router1#show ip arp 134.141.235.165
Protocol  Address          Age (min)  Hardware Addr Type   Interface
-------------------------------------------------------------------------
Internet  134.141.235.165   -       0002.1664.a5b3  ARPA   Vlan2
```

This example shows how to display an entry related to a VLAN:

```
Matrix>Router1#show ip arp vlan 2
Protocol  Address          Age (min)Hardware Addr Type   Interface
-------------------------------------------------------------------------
Internet  134.141.235.251   0      0003.4712.7a99 ARPA   Vlan2
```

Table 7-4 provides an explanation of the command output.

Table 7-4 show ip arp Output Details

Output	What It Displays...
Protocol	ARP entry's type of network address.
Address	Network address mapped to the entry's MAC address.
Age (min)	Interval (in minutes) since the entry was entered in the table.
Hardware Addr	MAC address mapped to the entry's network address.
Type	Encapsulation type used for the entry's network address.
Interface	Interface (VLAN or loopback) through which the entry was learned.

Adding or Removing Static ARP Entries

From global configuration mode, use this command to add or remove static (permanent) ARP table entries. Up to 1,000 static ARP entries are supported per Matrix Series routing module. A multicast MAC address can be used in a static ARP entry.

arp *ip-address mac-address* **arpa**

Use this "no" form of the command to remove a static ARP entry:

no arp *ip-address*

Disabling or Re-enabling Proxy ARP on an Interface

A variation of the ARP protocol, proxy ARP allows the routing module to send an ARP response on behalf of an end node to the requesting host. Proxy ARP can lessen bandwidth use on slow-speed WAN links. It is enabled by default.

From interface configuration mode, use this command to disable proxy ARP on an interface:

```
no ip proxy-arp
```

Use this form of the command to re-enable proxy ARP:

```
ip proxy-arp
```

Assigning a MAC Address to an Interface

By default, every routing interface uses the same MAC address. When necessary, you can assign different MAC address to different interfaces, as long as you are careful to select a MAC address that will not conflict with other devices on the VLAN.

From interface configuration mode, use this command to assign a MAC address on an interface.

```
ip mac-address address
```

Note: It is important to select a MAC address that will not conflict with other devices on the VLAN since the Matrix Series device will not automatically detect this conflict.

Use the “no” form of this command to remove a MAC address from an interface:

```
no ip mac-address address
```

Changing the ARP Timeout

The Matrix DFE Series or N-SA device can support up to 2000 outstanding unresolved ARP entries. By default, these entries will stay in the ARP table for 14,400 seconds before expiring.

From global configuration mode, use this command to change the ARP timeout value:

```
arp timeout seconds
```

Valid values are **0** - **65535** seconds. A value of **0** specifies that ARP entries will never be aged out.

Clearing the ARP Cache

From privileged EXEC mode, use this command to delete all nonstatic (dynamic) entries from the ARP table:

```
clear arp-cache
```

Configuring IP Broadcast Settings

Use the procedures in this section to perform the following IP broadcast configuration tasks:

- Enabling or Disabling IP Directed Broadcasts
- Configuring UDP Broadcast Forwarding

Enabling or Disabling IP Directed Broadcasts

Directed broadcasts are network or subnet broadcast packets which are sent to a router for forwarding. They can be misused to create Denial of Service (DoS) attacks. By default, the Matrix DFE Series or N-SA device protects against this possibility by not forwarding directed broadcasts. However, depending on your network requirements, you may want to enable this function.

From interface configuration mode, use this command to enable IP directed broadcasts on an interface:

```
ip directed-broadcast
```

Use this "no" form of the command to globally disable IP directed broadcasts:

```
no ip directed-broadcast
```

Configuring UDP Broadcast Forwarding

Typically, broadcast packets from one interface are not forwarded (routed) to another interface. However, some applications use UDP broadcasts to detect the availability of services, and protocols, such as BOOTP/DHCP, require broadcast forwarding to provide services to clients on other subnets. Configuring UDP broadcast forwarding on the Matrix DFE Series or N-SA device involves enabling it, enabling DHCP/BootP relay, and assigning an IP "helper address" as described in this section.

Enabling or Disabling UDP Broadcast Forwarding

From global configuration mode, use this command to enable UDP broadcast forwarding and to specify the protocols for which UDP will forward services:

```
ip forward-protocol {udp [port | protocol]}
```

The *port* variable specifies a destination port that controls which UDP services are forwarded. If not specified, the Matrix DFE Series or N-SA device will forward default services using the default ports listed in Table 7-5. Specifying a *protocol* keyword will enable only its service on the default port listed.

Table 7-5 Default UDP Forwarding Services

Protocol Keyword	Service	Default UDP Port
bootps	Bootstrap Protocol Server	67
bootpc	Bootstrap Protocol Client	68
domain	Domain Name Service	53
nameserver	EN-116 Name Service	42
netbios-dgm	NetBIOS datagram service	138
netbios-ns	NetBIOS name service	137
tacacs	Terminal Access Controller Access Control System	49
tftp	Trivial File Transfer Protocol	69
time	Time Service	37

Note: If a particular service exists inside the node, and there is no need to forward the request to remote networks, the “no” form of this command should be used to disable the forwarding for the specific port. Such requests will not be automatically blocked from being forwarded, just because a service for them exists in the node.

Use this “no” form of the command to remove a UDP port or protocol, disabling forwarding for the associated service:

```
no ip forward-protocol {udp [port | protocol]}
```

Enabling DHCP/BOOTP Relay

DHCP/BOOTP relay functionality is applied with the help of IP broadcast forwarding. A typical situation occurs when a host requests an IP address with no DHCP server located on that segment. A routing module can forward the DHCP request to a server located on another network if:

- IP forward-protocol is enabled for UDP as described in the previous section, and
- the address of the DHCP server is configured as a helper address on the receiving interface of the routing module forwarding the request, as described in the following section.

The DHCP/BOOTP relay function will detect the DHCP request and make the necessary changes to the header, replacing the destination address with the address of the server, and the source with its own address, and send it to the server. When the response comes from the server, the DHCP/BOOTP relay function sends it to the host.

From interface configuration mode, use this command to enable DHCP/BOOTP relay, and the forwarding of local UDP broadcasts to one or more specific destination addresses.

```
ip helper-address address
```

Example

This example shows how to permit to permit UDP broadcasts from hosts on networks 191.168.1.255 and 192.24.1.255 to reach servers on those networks:

```
Matrix>Router(config)#ip forward-protocol udp
Matrix>Router(config)#interface vlan 1
Matrix>Router(config-if(Vlan 1))#ip helper-address 192.168.1.255
Matrix>Router(config)#interface vlan 2
Matrix>Router(config-if(Vlan 2))#ip helper-address 192.24.1.255
```

Configuring Routes and Monitoring IP Traffic

Use the procedures in this section to perform the following IP route configuration and traffic monitoring tasks:

- Adding or Removing Static IP Routes
- Displaying IP Traffic and Performance
- Clearing IP Traffic Counters

Adding or Removing Static IP Routes

From global configuration mode, use this command to add a static IP route.

ip route *prefix mask* {*forward-addr* | **vlan** *vlan-id*} [*distance*] [**permanent**] [**tag** *value*]

The *forward-addr* or *vlan-id* specifies the next hop gateway.

Valid *distance* values are **1** - **255**. Routes with lower values receive higher preference in route selection. If not specified, the default value of 1 will be applied.

If **permanent** and **tag** are not specified, the route will be set as non-permanent with no tag assigned.

Examples

This example shows how to set IP address 10.1.2.3 as the next hop gateway to destination address 10.0.0.0. The route is assigned a tag of 1:

```
Matrix>Router1(config)#ip route 10.0.0.0 255.0.0.0 10.1.2.3 tag 1
```

This example shows how to set IP address 10.1.2.3 as the next hop gateway to destination address 10.0.0.0. The route is set as permanent and assigned a tag of 20:

```
Matrix>Router1(config)#ip route 10.0.0.0 255.0.0.0 10.1.2.3 permanent tag 20
```

This example shows how to set VLAN 100 as the next hop interface to destination address 10.0.0.0:

```
Matrix>Router1(config)#ip route 10.0.0.0 255.0.0.0 vlan 100
```

Use this “no” form of the command to remove an IP route:

no ip route *prefix mask* {*forward-addr* | **vlan** *vlan-id*}

Displaying IP Traffic and Performance

From any router mode, use the commands listed in Table 7-6 to display traffic statistics and other IP-related information.

Table 7-6 Show IP Traffic Commands

Command	Output
show ip traffic [**softpath**]	IP traffic statistics, including IP, ICMP, UDP, TCP, IGMP, and ARP traffic counters. The softpath option is used for debugging.
show ip route [*destination prefix destination prefix mask* **longer-prefixes | connected** | **ospf | rip | static** | **summary**]	Detailed or summary IP route information related to all known routes, or to connected routes, static routes, or those configured for the OSPF or RIP routing protocols.

Clearing IP Traffic Counters

From privileged EXEC mode, use this command to clear all IP traffic statistics counters (IP, ICMP, UDP, TCP, IGMP, and ARP).

```
clear ip stats
```

Configuring ICMP

Use the procedures in this section to perform the following ICMP configuration tasks:

- Enabling or Disabling ICMP
- Sending Ping Requests
- Using Traceroute

Enabling or Disabling ICMP

By default, ICMP (Internet Control Message Protocol) messaging is enabled on a Matrix DFE Series / N-SA routing interface for both echo-reply and mask-reply modes. If, for security reasons, ICMP has been disabled, use this command from interface configuration mode to re-enable ICMP. This will allow the router to reply to IP ping requests.

```
ip icmp {echo-reply | mask-reply}
```

Use this “no” form of the command to disable ICMP:

```
no ip icmp {echo-reply | mask-reply}
```

Sending Ping Requests

From privileged EXEC mode, use this command to test routing network connectivity by sending IP ping requests. The ping utility (IP ping only) transmits a maximum of five echo requests, with a packet size of 100. The application stops when the response has been received, or after the maximum number of requests has been sent.

```
ping ip-address
```

Examples

This example shows a successful ping to IP address 182.127.63.23:

```
Matrix>Router1#ping 182.127.63.23
Reply from 182.127.63.23
Reply from 182.127.63.23
Reply from 182.127.63.23

------ PING 182.127.63.23 : Statistics ------
 3 packets transmitted, 3 packets received, 0% packet loss
```

This example shows an unsuccessful ping to IP address 182.127.63.24:

```
Matrix>Router1#ping 182.127.63.24
Timed Out
Timed Out
Timed Out

 ------ PING 182.127.63.24 : Statistics ------
 3 packets transmitted, 0 packets received, 100% packet loss
```

Using Traceroute

From privileged EXEC mode, use this command to display a hop-by-hop path through an IP network from the device to a specific destination host. Three ICMP probes will be transmitted for each hop between the source and the traceroute destination.

traceroute *host*

Example

This example shows how to use traceroute to display a round trip path to host 192.167.252.46. In this case, hop 1 is an unnamed router at 192.167.201.2, hop 2 is “rtr10” at 192.4.9.10, hop 3 is “rtr43” at 192.167.208.43, and hop 4 is back to the host IP address. Round trip times for each of the three ICMP probes are displayed before each hop. Probe time outs are indicated by an asterisk (*):

```
Matrix>Router1#traceroute 192.167.225.46
Traceroute to 192.167.225.46, 30 hops max, 40 byte packets
1  10.00 ms  20.00 ms  20.00 ms  192.167.201.2 []
2  20.00 ms  20.00 ms  20.00 ms  192.4.9.10 [enatel-rtr10.enatel.com]
3  240.00 ms  *         480.00 ms  192.167.208.43 [enatel-rtr43.enatel.com]
4  <1 ms     *         20.00 ms  192.167.225.46 [enatel-rtr46.enatel.com]

TraceRoute Complete
```

8

Configuring OSPF

This chapter provides information about the following Open Shortest Path First (OSPF) protocol configuration procedures on the Matrix DFE Series or N-SA device.

For information about...	Refer to page...
Using OSPF on a Matrix DFE or N-SA Series Device	8-1
OSPF Overview	8-1
Configuring OSPF	8-7
Monitoring and Maintaining OSPF	8-12

*** Advanced License Required ***

OSPF is an advanced routing feature that must be enabled with a license key. If you have purchased an advanced routing license and have enabled routing on the device, you must activate your license as described in "Activating Advanced Routing" on page 8-9 in order to enable the OSPF command set. If you wish to purchase an advanced routing license, contact Enterasys Networks Sales.

Using OSPF on a Matrix DFE or N-SA Series Device

As discussed in "Using the Matrix DFE Series or N-SA Device as a Router" on page 7-1, operating the device in routing mode — as compared to switch mode only — allows users from one VLAN to communicate with another VLAN connected to the same device. When it becomes necessary for users from one network to communicate with users connected to another router on another network, a routing protocol such as OSPF must be configured on the routing interfaces. The procedures and sample configurations in this chapter describe how to configure OSPF on Matrix DFE Series of N-SA devices to serve that purpose.

OSPF Overview

Matrix DFE Series and N-SA devices support Open Shortest Path First (OSPF) Version 2.0, as defined in RFC 2328. OSPF is an interior gateway link state protocol that distributes routing information between routers in a single autonomous system (AS). OSPF chooses the least-cost path as the best path. It is a suitable routing protocol for complex networks with a large number of routers because it provides equal-cost multi-path routing, sending packets to a single destination by way of more than one interface simultaneously.

As an Interior Gateway Protocol (IGP), OSPF maintains a database of routing information it learns from other systems running the same protocol within the network or autonomous system (AS). Since OSPF is a link-state routing protocol, systems running OSPF flood the network when a change is detected, and maintain a link state database (LSDB). When there is a change, such as a

router being added to the network, OSPF uses the Shortest Path First (SPF) algorithm (also referred to as the Dijkstra algorithm) to calculate new routes.

OSPF Terminology

Table 8-1 defines key terminology used in OSPF configuration.

Table 8-1 OSPF Terms and Definitions

Term	Definition
ABR	Area Border Router located on the border of one or more OSPF area connecting those areas to the backbone network.
AS	Autonomous System — A collection of networks under a common administration sharing a common routing strategy
ASBR	Autonomous System Boundary Router — located between an OSPF Autonomous System and a non-OSPF network
ASE	Autonomous System External destination
LSA	Link State Advertisement — Broadcast packet containing information about neighbors and path costs used by receiving routers to maintain routing tables
LSDB	Link State Database
NSSA	Not-So-Stubby Area — A form of stub area that allows some advertisements from external routes, used when an OSPF network is connected to multiple non-OSPF routing domains
Stub area	Area that carries a default route, intra-area routes, and inter-area routes, but does not carry external routes
Virtual link	Logical connection between an OSPF backbone and non-backbone area.

Supported Functions

Matrix DFE Series and N-SA devices support the following OSPF functions:

- Authentication: Simple password and MD5 authentication methods are supported within an area.
- One OSPF Instance per module: Used to enable or disable OSPF.
- Interface Parameters: Parameters that can be configured include interface output cost, retransmission interval, interface transit delay, router priority, router dead and hello intervals, and authentication.
- RIP Route Redistribution: Routes learned by way of the RIP routing protocol, as well as static and connected routes, can be redistributed into OSPF. OSPF routes can also be redistributed into RIP.
- Definition of stub areas
- Definition of Not So Stubby Areas (NSSAs)
- Static multi-path forwarding
- Virtual links

Link-State Advertisements (LSAs)

Using OSPF, the Matrix DFE-Platinum Series system floods link-state advertisement (LSA) packets to keep its topological database updated, and to help ensure the databases of its neighbors are also current.

Types of LSAs, each relating to a particular part of the OSPF routing domain, are as follows.

Router link advertisements are sent by each DFE Series system configured with OSPF to describe the router's links within the area. The router floods these advertisements within the single area only.

Network link advertisements describe all the routing systems that are attached to the network. The advertisements are flooded within a single area only. They are generated by the designated router.

Summary link advertisements describe the routes in other OSPF areas that the DFE Series system knows about. The DFE Series system must be configured as an area border router (ABR) to perform this type of flooding. These advertisements are flooded to the areas that use the routes. The areas must be within the same AS.

AS external link advertisements describe external routes that the DFE Series system knows about. The DFE Series system must be configured as an AS boundary router (ASBR) to flood this type of advertisement. These advertisements are flooded throughout the AS, except to stub areas. The DFE Series / N-SA system supports opaque Type 7 LSAs for not-so-stubby areas (NSSAs).

Matrix DFE Series and N-SA Operation in OSPF Networks

Each routing system within an OSPF AS maintains a database of the network topology. To reduce the size of the database, an AS is typically divided into smaller routing areas, and the routing systems perform different roles within the AS. You can configure the Matrix DFE-Platinum Series system to perform any of the three routing roles within an OSPF autonomous system: internal router, area border router (ABR), or AS boundary router (ASBR).

Figure 8-1 shows a sample OSPF topology, with the different types of routers.

Internal router – Systems within a specific area are called internal routers. When the TN250G functions as an internal router, all interfaces are directly connected to other routers within the same area.

Area border router – As an ABR, the DFE Series system can belong to multiple areas. In this role, the router maintains a separate topology database for each area in which it is a member. The DFE Series system's ABR support includes route summarization, which enables the router to advertise a route that consolidates multiple advertised addresses. The DFE Series system also supports virtual links, enabling connection to a backbone without a physical connection. Stub area support includes the ability to prevent ABRs from sending summary link-state advertisements into the stub area.

If the DFE Series system has more than one area configured, and at least one is the backbone, the router automatically is an ABR.

AS boundary router – The DFE Series system can also be configured as an ASBR. In this role, the router exchanges externally learned routing information with other routers throughout the AS. As an ASBR, the router provides route redistribution, as well as route aggregation to help reduce the amount of routing information in the network. The importing and exporting of routing policies is also a function of ASBRs. ASBRs are created when you configure route redistribution.

Figure 8-1 OSPF Topology

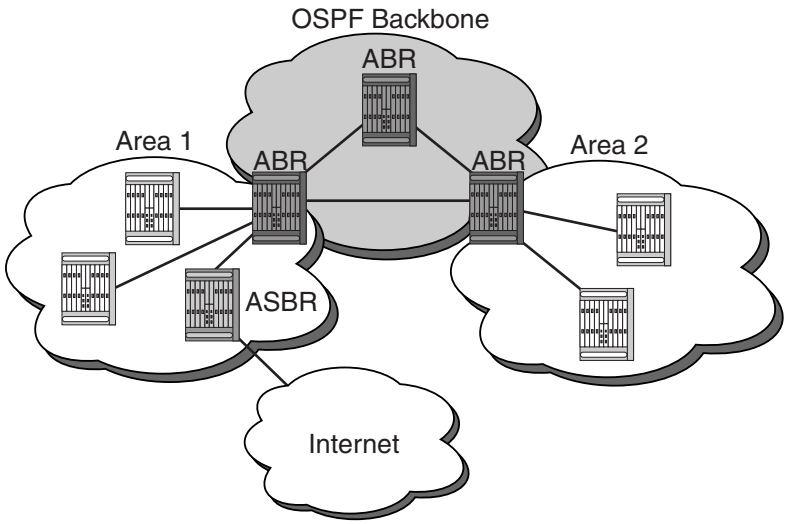

OSPF Areas

OSPF allows networks to be grouped into areas — a collection of subnets that are grouped in a logical fashion. These areas communicate with other areas through the backbone area. Routing information passed between areas is abstracted, potentially allowing a significant reduction in routing traffic. OSPF uses four different types of routes, listed in order of preference:

- Intra-area — Destinations within the same area
- Inter-area — Destinations within different areas
- Type 1 ASE — Autonomous System External destinations
- Type 2 ASE — Autonomous System External destinations

Both types of ASE routes are routes to destinations external to OSPF (and usually external to the AS). Routes exported into OSPF ASE as type 1 ASE routes are supposed to be from interior gateway protocols, such as RIP, whose external metrics are directly comparable to OSPF metrics. When a routing decision is being made, OSPF will add the internal cost to the AS border router to the external metric. Type 2 ASEs are used for exterior gateway protocols whose metrics are not comparable to OSPF metrics. In this case, only the external OSPF cost from the AS border router is used in the routing decision.

All imported routes default to external type 2 routes, and can be configured as external type 1 in cases where the metrics are comparable.

To reduce the amount of routing information propagated between areas, you can configure summary ranges on Area Border Routers (ABRs). Intra-area Link State Advertisements (LSAs)

that fall within the specified ranges are not advertised into other areas as inter-area routes. Instead, the specified ranges are advertised as summary network LSAs.

Note: Although this does not apply to most changes to OSPF and other routing-based entries in the configuration file, the following actions force the OSPF Link State Databases (LSDB) to re-initialize:

- Adding a network to or removing one from an area.
- Changing an area's type.
- Adding a summary range to or removing one from an Area Border Router.

Matrix DFE Series and N-SA devices allow configuration of various OSPF area parameters, including stub areas, default cost, authentication method, and not-so-stubby areas (NSSAs). Information about routes which are external to the OSPF routing domain is not sent into a stub area. Instead, there is a default external route generated by the ABR into the stub area for destinations outside the OSPF routing domain.

Stub cost specifies the cost to be used to inject a default route into a stub area. An authentication method for OSPF packets can be specified on a per-area basis.

Static Multipath Forwarding

The Matrix DFE-Platinum Series system supports OSPF and static multi-path. If multiple, equal-cost OSPF or static routes have been defined for any destination, then the DFE Series system "discovers" and uses all of them. The system will automatically learn up to four equal-cost OSPF or static routes and retain them in its Forwarding Information Base (FIB). The forwarding module then installs flows for these destinations in a round-robin fashion, or hash-based algorithm.

Designating an OSPF Instance

OSPF is enabled by creating an OSPF instance, which is designated by an identification number from 1 to 65535.

Note: Only one OSPF instance is allowed per Matrix DFE Series routing module or N-SA device.

Once an instance is created, all the router's OSPF settings are configured with respect to the instance ID.

Virtual Links

In OSPF, virtual links can be established:

- To connect an area via a transit area to the backbone
- To create a redundant backbone connection through another area

Each Area Border Router must be configured with the same virtual link. Note that virtual links cannot be configured through a stub area.

Autonomous System External (ASE) Link Advertisements

An autonomous system boundary router (ASBR) advertises external destinations throughout the OSPF autonomous system. In many cases, external link states make up a large percentage of the link states in the databases of every router. A stub area is an area in which you do not allow advertisements of external routes, thus reducing the size of the database even more. Instead, a default summary route (0.0.0.0) is inserted into the stub area in order to reach these external routes. Table 8-2 lists the restrictions associated with each type of area:

Table 8-2 OSPF Area Restrictions

Area	Restrictions
Normal	None
Stub area	No Type 5 AS-external LSA allowed
Totally Stub area	No Type 3, 4 or 5 LSAs allowed except the default summary route
NSSA	No Type 5 AS-external LSAs allowed, but Type 7 LSAs that convert to Type 5 at the NSSA ABR can traverse
NSSA Totally Stub area	No Type 3, 4 or 5 LSAs except the default summary route, but Type 7 LSAs that convert to Type 5 at the NSSA ABR are allowed

Router Modes Used for OSPF Configuration

As shown back in Table 7-2 on page 7-3, the Matrix DFE Series/N-SA CLI provides different modes of router operation for issuing a subset of commands from each mode. Table 8-3 lists the modes of operation used in configuring OSPF. Sample configurations provided in this chapter will demonstrate enabling these command modes.

Notes: Sample prompts shown in Table 8-3 assume module 1(or the N-SA device) has been configured for routing, and that VLAN 1 will be designated as a outing interface.

A *module* designation of 1 must be entered to enable routing on the Matrix N-SA standalone device. All other values will result in an error message.

Table 8-3 Router Modes Used for OSPF Configuration

Use this mode...	To...	Access method...	Resulting Prompt...
Privileged EXEC Mode	Set system operating parameters Show configuration parameters Save/copy configurations	From the switch CLI: 1. Type **router** *module* (using a module number configured for routing). 2. Type **enable**.	`Matrix>Router1>` `Matrix>Router1#`
Global Configuration Mode	Set system-wide parameters.	Type **configure terminal** from Privileged EXEC mode.	`Matrix>Router1(config)#`

Table 8-3 Router Modes Used for OSPF Configuration (continued)

Use this mode...	To...	Access method...	Resulting Prompt...
Interface Configuration Mode	Configure router interfaces.	Type **interface vlan** or **interface loopback** and the interface's *id* from Global Configuration mode.	`Matrix>Router1 (config-if(Vlan 1 | Lpbk 1))#`
OSPF Router Configuration Mode	Set OSPF parameters.	Type **router ospf** and the OSPF *process-id* from Global or Interface Configuration mode.	`Matrix>Router1 (config-router)#`

Note: To jump to a lower configuration mode, type **exit** at the command prompt. To revert back to switch CLI, type **exit** from Privileged EXEC router mode.

Configuring OSPF

Use the procedures in this section to perform the following OSPF configuration tasks:

- Modifying Default Configuration Settings
- Activating Advanced Routing
- Performing a Basic OSPF Configuration
- Configuring an OSPF NSSA

Modifying Default Configuration Settings

By default, OSPF is disabled on the device, and no instances or areas are configured. Once an OSPF instance is created and a network area assigned, the Matrix DFE Series or N-SA system will apply default settings for cost, timers, and other settings. Table 8-4 lists OSPF parameters, their default settings, and the associated commands and configuration modes used to modify defaults.

Table 8-4 OSPF Default Settings

Parameter	Default Setting	Config Mode	Command
ABR address ranges	None configured	Router (OSPF)	**area** *area-id* **range** *ip-address ip-mask*
Advanced routing license	None activated	Global	**license advanced** *license-key*
Areas	None configured	Global	**network** *ip-address wildcard-mask* **area** *area-id*
Authentication (area)	Disabled	Router (OSPF)	**area** *area-id* **authentication** {**simple** | **message-digest**}
Authentication (MD5)	None configured	Interface	**ip ospf message-digest-key** *keyid* **md5** *key*

Table 8-4 OSPF Default Settings (continued)

Parameter	Default Setting	Config Mode	Command
Authentication (simple text)	None configured	Interface	**ip ospf authentication-key** *password*
Cost (interface)	10	Interface	**ip ospf cost** *cost*
Cost (stub area from ABR)	1	Router (OSPF)	**area** *area-id* **default-cost** *cost*
Dead interval	40 seconds	Interface	**ip ospf dead-interval** *seconds*
Database overflow	Not configured	Router (OSPF)	**database-overflow external** {[**exit-overflow-interval** *interval*] [**limit** *limit*] [**warning-level** *level*]}
Distance	110	Router (OSPF)	**distance ospf** {**external** | **inter-area** | **intra-area**} *weight*
Enabled/disabled state	Disabled	Global	**router ospf** *process-id*
Hello interval	10 seconds (broadcast and PPP) 30 seconds (non-broadcast and point-to-multipoint networks)	Interface	**ip ospf hello-interval** *seconds*
Instance ID	None configured	Global	**router ospf** *process-id*
NSSA area	None configured	Router (OSPF)	**area** *area-id* **nssa** [**default-information-originate**]
Passive interface	None configured	Router (OSPF)	**passive-interface vlan** *vlan-id*
Priority	1	Interface	**ip ospf priority** *number*
Redistribution	Not configured	Router (OSPF)	**redistribute** {**rip** | **static** [**metric** *metric value*] [**metric-type** *type-value*] [**subnets**] [**tag**] | {**connected** [**route-map** *id-namber*] [**metric** *metric value*] [**metric-type** *type-value*] [**subnets**] [**tag** *tag*]}
Retransmit interval (LSAs)	1 second	Interface	**ip ospf retransmit-interval** *seconds*
Router ID	None configured	Global	**router ospf** *process-id*
Stub area	None configured	Router (OSPF)	**area** *area-id* **stub** [**no-summary**]
Transmit delay (LSAs)	1 second	Interface	**ip ospf transmit-delay** *seconds*
Virtual link	None configured	Router (OSPF)	**area** *area-id* **virtual-link** *ip-address* [**authentication-key** *key*][**dead-interval** | **hello-interval** | **retransmit-interval** | **transmit-delay** {*seconds*}]

Activating Advanced Routing

In order to enable advanced routing protocols, such as OSPF and extended ACLs, on a Matrix DFE Series or N-SA device, you must purchase and activate a license key. If you have purchased an advanced routing license, and have enabled routing on the device as described back in Procedure 3-1, you can proceed to activate your license as described in this section. If you wish to purchase an advanced routing license, contact Enterasys Networks Sales.

1. Determine your appropriate *license-key* from the activation key document shipped with your product, or by viewing the top of the **running-config** output:

   ```
   show running-config
   ```

2. Enter the applicable license key for advanced routing:

   ```
   set license advanced license-key [slot slot]
   ```

3. Activate the license by resetting the device:

   ```
   reset [mod-num | system ]
   ```

4. Verify license activation:

   ```
   show license
   ```

Performing a Basic OSPF Configuration

As described previously, while a basic router configuration allows users in VLANs attached to the same network device to communicate, traffic exchange between two or more networks requires the configuration of a routing protocol, such as OSPF.

With the Matrix DFE Series / N-SA advanced routing license activated, OSPF is enabled by first creating an OSPF instance, and then associating one or more network areas for that instance. Procedure 8-1 lists the steps and the associated commands necessary to complete this process. Once completed, optional configuration tasks can be performed, including changing the default interface and area parameters listed in Table 8-4.

Procedure 8-1 assumes that, as shown in the topology in Figure 8-2, the administrator wants to allow VLAN 100 users from Router 1 to communicate with users on neighboring Router 2. It assumes that advanced routing has been enabled, and that the administrator has configured VLAN 100 as a routing interface as described back in the "Sample Configuration" on page 7-8. In this case, the administrator has assigned the interface IP address 131.108.1.1 within OSPF area 1. Once the administrator has connected this interface from the port assigned to VLAN 100 on Router 1 over to Router 2, a link will be established with an IP address in the same network as VLAN 100. Because VLAN 100 frames will be untagged, there is no relationship between VLAN 100 on Router 1, and any VLAN on the neighboring router. Router 2 may not even have a VLAN 100. It will, nonetheless, believe it is a member of network 131.108.1.1 along with Router 1.

For a description of all the router configuration modes used in the Matrix DFE Series / N-SA CLI, refer back to Table 3-6.

Figure 8-2 Topology for Basic OSPF Configuration

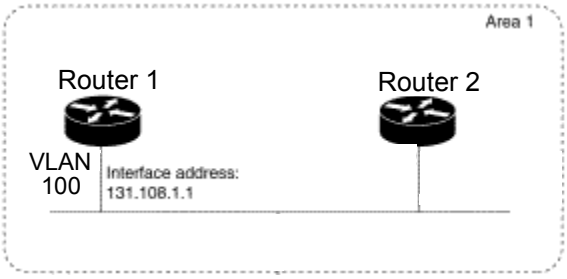

Procedure 8-1 Basic OSPF Configuration

Step	Task	Command
1.	From switch mode, enter router mode on router 1.	**router 1**
2.	Enable router 1.	**enable**
3.	Enable global configuration mode.	**configure terminal**
4.	Create routing instance 1 on router 1, and enable OSPF router configuration mode for this instance.	**router ospf 1**
5.	Assign network interface 131.108.1.1 to area 1.	**network 131.108.1.1 255.255.255.0 area 1**

Example

This example shows the complete configuration described in Procedure 8-1, beginning with VLAN 100 configuration in switch mode.

```
**Create the VLAN**
Matrix(rw)->set vlan create 100

**Set the inbound PVID associations for untagged frames**
Matrix(rw)->set port vlan fe.5.21-40 100

**Set the list of egress ports and tagging for the VLAN**
Matrix(rw)->set vlan egress 100 fe.5.21-20 untagged

**Configure the router**
Matrix(rw)->set router 1
Matrix(rw)->router 1
Matrix>Router1>enable
Matrix>Router1#configure terminal

**Assign and enable the routing interface**
Matrix>Router1(config)#interface vlan 100
Matrix>Router1(config-if(Vlan 100))#ip address 131.108.1.1 255.255.255.0
Matrix>Router1(config-if(Vlan 100))#no shutdown
Matrix>Router1(config-if(Vlan 100))#exit
**Configure OSPF on the interface**
Matrix>Router1(config)#router ospf 1
Matrix>Router1(config-router)#network 131.108.1.1 255.255.255.0 area 1
```

Configuring an OSPF NSSA

As described in "OSPF Areas" on page 8-4, OSPF networks are grouped into areas. Routing information passed between areas is abstracted, potentially allowing a significant reduction in routing traffic.

An OSPF area can be configured as:

- a normal area (default), which would not restrict Link State Advertisements (LSAs),
- a stub area, which would not allow LSAs from external routes, or
- a Not So Stubby Area (NSSA), which allows some advertisements from external routes.

While the stub area reduces the size of the link state database (LSDB), NSSA area configuration allows LSDB reduction while still permitting the importing of selected external routes. External routes that are not imported into an NSSA can be represented by means of a default route. This configuration is used when an OSPF internetwork is connected to multiple non-OSPF routing domains.

Procedure 8-2 lists the steps and the associated commands necessary to complete the NSSA configuration process. It assumes the same topology used in Performing a Basic OSPF Configuration on page 8-9, only this example configures area 1 as an NSSA, adds security parameters, and establishes a virtual link between area 1 and the OSPF backbone.

Procedure 8-2 OSPF NSSA Configuration

Step	Task	Command
1.	From switch mode, enter router mode on router 1.	`router 1`
2.	Enable router 1.	`enable`
3.	Enable global configuration mode.	`configure terminal`
4.	Enable interface configuration mode for VLAN 100.	`interface vlan 100`
5.	Configure MD5 authentication key 100, and the associated password "ospfkey" to be used by neighboring routers.	`ip ospf message-digest-key 100 md5 ospfkey`
6.	Enable router configuration mode for OSPF instance 1.	`router ospf 1`
7.	Assign network interface 131.108.1.1 to area 1.	`network 131.108.1.1 255.255.255.0 area 1`
8.	Configure the area 1 router as an ABR, and set the scope of the area.	`area 1 range 131.108.1.1 255.255.255.0`
9.	Enable MD5 authentication for area 1.	`area 1 authentication message-digest`
10.	Configure area 1 as an NSSA, allowed to generate Type 7 LSAs.	`area 1 nssa`
11.	Create a virtual link with MD5 authentication for the ABR to represent a logical connection between area 1 and the OSPF backbone.	`area 1 virtual-link 131.108.1.3 authentication-key 100`

Example

This example shows the complete configuration described in Procedure 8-2, beginning with VLAN 100 configuration in switch mode.

```
**Create the VLAN**
Matrix(rw)->set vlan create 100

**Set the inbound PVID associations for untagged frames**
Matrix(rw)->set port vlan fe.5.21-40 100

**Set the list of egress ports and tagging for the VLAN**
Matrix(rw)->set vlan egress 100 fe.5.21-20 untagged

**Configure the router**
Matrix(rw)->set router 1
Matrix(rw)->router 1
Matrix>Router1>enable
Matrix>Router1#configure terminal

**Assign and enable the routing interface**
Matrix>Router1(config)#interface vlan 100
Matrix>Router1(config-if(Vlan 100))#ip address 131.108.1.1 255.255.255.0
Matrix>Router1(config-if(Vlan 100))#no shutdown


**Configure MD5 authentication on the interface**
Matrix>Router1(config-if(Vlan 100))#ip ospf message-digest-key 100 md5 ospfkey
Matrix>Router1(config-if(Vlan 100))#exit

**Configure OSPF area 1 as an NSSA**
Matrix>Router1(config)#router ospf 1
Matrix>Router1(config-router)#network 131.108.1.1 255.255.255.0 area 1
Matrix>Router1(config-router)#area 1 range 131.108.1.1 255.255.255.0
Matrix>Router1(config-router)#area 1 authentication message-digest
Matrix>Router1(config-router)#area 1 nssa
Matrix>Router1(config-router)#area 1 virtual-link 131.108.1.3 authentication-key
   100
```

Monitoring and Maintaining OSPF

Use the procedures in this section to monitor and maintain OSPF on your system:

- Displaying OSPF Information
- Displaying OSPF Information
- Resetting OSPF
- Debugging OSPF

Displaying OSPF Information

From any router mode, use the commands listed in Table 8-5 to display OSPF information.

Table 8-5 Displaying OSPF Information

Command	Output
show ip ospf	OSPF instance information, including area, interfaces, and global parameters.
show ip ospf border-routers	Information about Autonomous System (AS) and Area Border Routers (ABRs).
show ip ospf database [**database-summary**] | [**router** | **network** | **summary** | **asbr-summary** | **external** | **nssa-external** [*link-state-id*]	One or all link state database records, a numerical summary of the LSDB contents, or only record(s) related to a specific link state type and/or link state ID.
show ip ospf interface [**vlan** *vlan-id*]	Interface related information, including network type, priority, cost, hello interval, and dead interval, for one or all routing VLANs.
show ip ospf neighbor [**detail**][*ip-address*] [**vlan** *vlan-id*]	Detailed or summary information about OSPF routing neighbors assigned to one or all IP addresses and/or VLANs.
show ip ospf virtual-links	All OSPF virtual links.

Limiting Database Overflow

OSPF link state database overflow occurs when the router is unable to maintain the database in its entirety due to excessive external LSAs. Setting database overflow allows you to limit the number of external LSAs. If the limit is exceeded, self-originated external LSAs will be removed so that OSPF can handle the large number of external LSAs coming from another router. When the warning level is set, a Syslog message will be issued when the number of external LSAs has reached the specified level. After the specified **exit-overflow-interval** is reached, the database will be checked and, if the LSA total is less than the limit specified, the self originated external LSAs will be restored.

From OSPF configuration mode, use the following command to limit OSPF database overflow:

database-overflow external {[**exit-overflow-interval** *interval*] [**limit** *limit*] [**warning-level** *level*]}

Exit-overflow-interval specifies an interval (in seconds) the OSPF link state database will be checked to determine if the overflow limit has been reached. Valid values are **0** - **86400**. Default is **0**.

Limit specifies the peak number of LSAs accepted before overflow occurs. Valid values are **0** - **4000**. Default is **0**. Limit value must be greater than the **warning-level** value and set prior to it since all defaults are 0.

Warning-level specifies the number of LSAs at which a warning of pending overflow will be generated. Valid values are **0** - **4000**. Default is **0**.

Resetting OSPF

From Privileged EXEC mode, use this command to reset the OSPF process, forcing adjacencies to be reestablished and routes to be reconverged.

clear ip ospf process *process-id*

Debugging OSPF

From Privileged EXEC mode, use this command to generate OSPF debugging output displaying adjacency, flooding, retransmission events, LSA generation or packet processing information.

debug ip ospf {**adj** | **flood** | **lsa-generation** | **packet** | **retransmission**}

9

Configuring VRRP

This chapter provides the following information about configuring and monitoring the Virtual Router Redundancy Protocol (VRRP) on the Matrix DFE Series or N-SA device:

For information about...	Refer to page...
VRRP Overview	9-1
Configuring VRRP	9-2
Modifying a Configuration	9-13
Monitoring VRRP	9-17
VRRP Configuration Notes	9-17

VRRP Overview

This section provides a brief overview of VRRP configuration on the Matrix DFE Series or N-SA device. VRRP is defined in RFC 3678. This protocol eliminates the single point of failure inherent in the static default routed environment by transferring the responsibility from one router to another if the original router goes down.

End host systems on a LAN are often configured to send packets to a statically configured default router. If this default router becomes unavailable, all the hosts that use it as their first hop router become isolated on the network. VRRP provides a way to ensure the availability of an end host's default router.

This is done by assigning IP addresses that end hosts use as their default route to a "virtual router." A master router is assigned to forward traffic designated for the virtual router. If the master router should become unavailable, a backup router takes over and begins forwarding traffic for the virtual router. As long as one of the routers in a VRRP configuration is up, the IP addresses assigned to the virtual router are always available, and the end hosts can send packets to these IP addresses without interruption.

VRRP Terms and Definitions

Table 9-1 lists terms and definitions used in VRRP configuration.

Table 9-1 VRRP Terms and Definitions

Term	Definition
VRID	Virtual Router ID — a unique number associated with each VRRP routing interface.

Table 9-1 VRRP Terms and Definitions (continued)

Term	Definition
master	The router assigned to forward traffic designated for the virtual router. The master sends an advertisement to all other VRRP routers declaring its status and assumes responsibility for forwarding packets associated with its VRID.
backup	The router that takes over and begins forwarding traffic for the virtual router if the master router becomes unavailable. When more than one backup is configured for a VRID, then the priority value (as described below) will determine whether a router will act as a primary or secondary backup.
virtual router	Designates an IP address and VRID associated with a default route to use if the master router becomes unavailable. As long as one of the routers in a VRRP configuration is up, the IP addresses assigned to the virtual router are always available, and the end hosts can send packets to these IP addresses without interruption.
owner	A value used in setting a VRRP address which indicates if a router owns an IP address. A value or 1 indicates the router owns the address and is, therefore, the master router for that interface. A value of 0 indicates the router does not own the address and is a backup for the interface.
priority	Specifies a VRRP priority value to associate with a VRID. Valid values are from 1 to 254, with the highest value setting the highest priority. If the virtual router IP address is not owned by any of the VRRP routers, then the routers compare their priorities and the higher priority owner becomes the master. If priority values are the same, then the VRRP router with the higher IP address is selected master. **Note:** Priority value of 255 is reserved for the VRRP router that owns the IP address associated with the virtual router. Priority 0 is reserved for signaling that the master has stopped working and the backup router must transition to master state.

Configuring VRRP

Important Notice

The configuration examples provided in this section assume that IP routing has been enabled on all routing modules as described in Chapter 7, Configuring IP. Examples use VLAN 1 as the configured routing interface. Depending on your configuration, your router and VLAN IDs may be different.

This section presents the following sample VRRP configurations:

- A basic VRRP configuration with one virtual router (page 9-3)
- A symmetrical VRRP configuration with two virtual routers (page 9-5)
- A multi-backup VRRP configuration with three virtual routers (page 9-8)

Notes: In some instances, a heavy load on a VRRP master may delay VRRP packet transmission and cause the backup router to assume the role of master.

Each Matrix DFE Series routing module or N-SA device supports up to 64 VRRP sessions. A session is defined as one virtual router running on one interface. Running a single virtual router on four interfaces is considered four sessions of VRRP, as is running four virtual routers on a single interface.

Do not use an IP address for VRRP that is already configured for load-balancing.

Basic VRRP Configuration

Figure 9-1 shows a basic VRRP configuration with a single virtual router. Routers R1 and R2 are both configured with one virtual router (VRID 1). Router R1 serves as the master and Router R2 serves as the backup. The four end hosts are configured to use 10.0.0.1/16 as the default route. IP address 10.0.0.1/16 is associated with virtual router ID (VRID) 1.

Figure 9-1 Basic VRRP Configuration

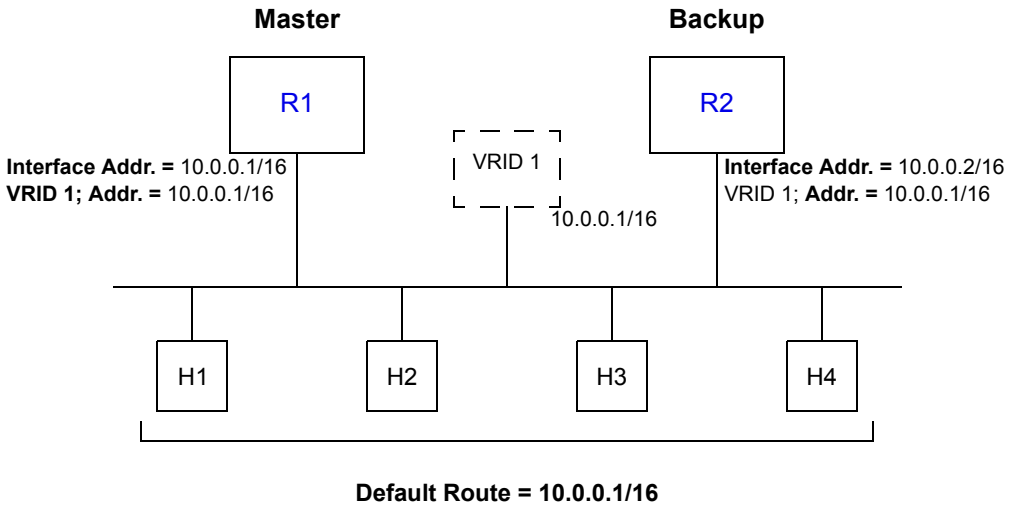

If Router R1 should become unavailable, Router R2 would take over virtual router VRID 1 and its associated IP addresses. Packets sent to 10.0.0.1/16 would go to Router R2. When Router R1 comes up again, it would take over as master, and Router R2 would revert to backup.

Router R1 Configuration

Use the following procedure to configure Router R1 as shown back in Figure 9-1:

Procedure 9-1 Configuring Router 1 for Basic VRRP

Step	Task	Command(s)
1.	In **Global Configuration mode**: Create a routing interface for Router 1 on VLAN 1.	`interface vlan 1`
2.	In **Interface Configuration mode**: Set and enable IP address 10.0.0.1 255.255.255.0 on this router and VLAN.	`ip address 10.0.0.1 255.255.255.0` `no shutdown`
3.	In **Global Configuration mode**: Enable VRRP configuration mode on this router.	`router vrrp`

Procedure 9-1 Configuring Router 1 for Basic VRRP (continued)

Step	Task	Command(s)
4.	In **VRRP Configuration mode**: Create a VRRP session for this router on VLAN 1 with a VRID of 1.	`create vlan 1 1`
5.	Set a virtual router address as VLAN 1, VRID 1, interface 10.0.0.1, and set this router as the master (owner value 1).	`address vlan 1 1 10.0.0.1 1`
6.	Enable VRRP on this router on VLAN 1, VRID 1. **Note:** Before enabling VRRP on this router you must set the other options described in this section. Once enabled, you cannot make any configuration changes to VRRP on this router without first disabling VRRP using the **no enable vlan** command.	`enable vlan 1 1`

Router R2 Configuration

Use the following procedure to configure Router R2 as shown back in Figure 9-1. Notice that the configuration for Router R2 is nearly identical to Router R1. The difference is that Router R2 does not own IP address 10.0.0.1/16. Since Router R2 does not own this IP address, it is the backup and will take over from the master if it should become unavailable.

Procedure 9-2 Configuring Router 2 for Basic VRRP

Step	Task	Command(s)
1.	In **Global Configuration** mode: 1. Create a routing interface for Router 2 on VLAN 1.	`interface vlan 1`
2.	In **Interface Configuration** mode: Set and enable IP address 10.0.0.2 255.255.255.0 on this router and VLAN.	`ip address 10.0.0.2 255.255.255.0` `no shutdown`
3.	In **Global Configuration** mode: Enable VRRP configuration mode on this router.	`router vrrp`
4.	In **VRRP Configuration** mode: Create a VRRP session for this router on VLAN 1 with a VRID of 1.	`create vlan 1 1`
5.	Set a virtual router address as VLAN 1, VRID 1, interface 10.0.0.1, and set this router as backup (owner value 0).	`address vlan 1 1 10.0.0.1 0`
6.	Enable VRRP on this router on VLAN 1, VRID 1. **Note:** Before enabling VRRP on this router you must set the other options described in this section. Once enabled, you cannot make any configuration changes to VRRP on this router without first disabling VRRP using the **no enable vlan** command.	`enable vlan 1 1`

Example: Basic VRRP Configuration

This example shows the complete configuration for basic VRRP on Routers 1 and 2:

```
**Router 1**
Matrix>Router1(config)#interface vlan 1
Matrix>Router1(config-if(Vlan 1))#ip address 10.0.0.1 255.255.255.0
Matrix>Router1(config-if(Vlan 1))#no shutdown
Matrix>Router1(config)#router vrrp
Matrix>Router1(config-router)#create vlan 1 1
Matrix>Router1(config-router)#address vlan 1 1 10.0.0.1 1
Matrix>Router1(config-router)#enable vlan 1 1

**Router 2**
Matrix>Router2(config)#interface vlan 1
Matrix>Router2(config-if(Vlan 1))#ip address 10.0.0.2 255.255.255.0
Matrix>Router2(config-if(Vlan 1))#no shutdown
Matrix>Router2(config)#router vrrp
Matrix>Router2(config-router)#create vlan 1 1
Matrix>Router2(config-router)#address vlan 1 1 10.0.0.1 0
Matrix>Router2(config-router)#enable vlan 1 1
```

Symmetrical Configuration

Figure 9-2 shows a VRRP configuration with two routers and two virtual routers. Routers R1 and R2 are both configured with two virtual routers (VRID 1 and VRID 2).

Router R1 serves as:

- Master for VRID 1
- Backup for VRID 2

Router R2 serves as:

- Master for VRID 2
- Backup for VRID 1

This configuration allows you to load-balance traffic coming from the hosts on the 10.0.0.0/16 subnet and provides a redundant path to either virtual router.

Note: This is the recommended configuration on a network using VRRP.

Figure 9-2 Symmetrical VRRP Configuration

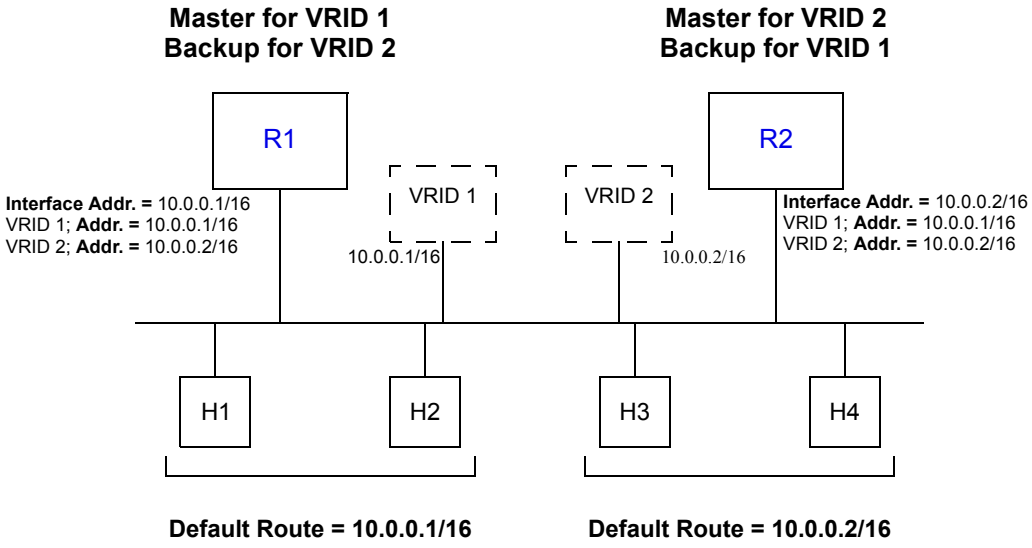

In this configuration, half the hosts use 10.0.0.1/16 as their default route, and half use 10.0.0.2/16. IP address 10.0.0.1/16 is associated with virtual router VRID 1, and IP address 10.0.0.2/16 is associated with virtual router VRID 2.

If Router R1, the master for virtual router VRID 1, goes down, Router R2 would take over the IP address 10.0.0.1/16. Similarly, if Router R2, the master for virtual router VRID 2, goes down, Router R1 would take over the IP address 10.0.0.2/16.

Router R1 Configuration

Use the following procedure to configure Router R1 as shown back in Figure 9-2:

Procedure 9-3 Configuring Router 1 for Symmetrical VRRP

Step	Task	Command(s)
1.	In **Global Configuration** mode: Create a routing interface for Router 1 on VLAN 1.	`interface vlan 1`
2.	In **Interface Configuration** mode: Set and enable IP address 10.0.0.1 255.255.255.0 on this router and VLAN.	`ip address 10.0.0.1 255.255.255.0` `no shutdown`
3.	In **Global Configuration** mode: Enable VRRP configuration mode on this router.	`router vrrp`
4.	In **VRRP Configuration** mode: Create a VRRP session for this router on VLAN 1 with a VRID of 1.	`create vlan 1 1`
5.	Create a second VRRP session on VLAN 1 with a VRID of 2.	`create vlan 1 2`

Procedure 9-3 Configuring Router 1 for Symmetrical VRRP (continued)

Step	Task	Command(s)
6.	Set a virtual router address as VLAN 1, VRID 1, interface 10.0.0.1, and set this router as the master (owner value 1).	`address vlan 1 1 10.0.0.1 1`
7.	Set a virtual router address as VLAN 1, VRID 2, interface 10.0.0.2, and set this router as backup (owner value 0).	`address vlan 1 2 10.0.0.2 0`
8.	Enable VRRP on this router on VLAN 1, VRIDs 1 and 2.	`enable vlan 1 1` `enable vlan 1 2`

Note: Before enabling VRRP on this router you must set the other options described in this section. Once enabled, you cannot make any configuration changes to VRRP on this router without first disabling VRRP using the **no enable vlan** command.

Router R2 Configuration

Use the following procedure to configure Router R2 as shown in Figure 9-2:

Procedure 9-4 Configuring Router 2 for Symmetrical VRRP

Step	Task	Command(s)
1.	In **Global Configuration** mode: Create a routing interface for Router 2 on VLAN 1.	`interface vlan 1`
2.	In **Interface Configuration** mode: Set and enable IP address 10.0.0.2 255.255.255.0 on this router and VLAN.	`ip address 10.0.0.2 255.255.255.0` `no shutdown`
3.	In **Global Configuration** mode: Enable VRRP configuration mode on this router.	`router vrrp`
4.	In **VRRP Configuration** mode: Create a VRRP session for this router on VLAN 1 with a VRID of 1.	`create vlan 1 1`
5.	Create a second VRRP session on VLAN 1 with a VRID of 2.	`create vlan 1 2`
6.	Set a virtual router address as VLAN 1, VRID 2, interface 10.0.0.2, and set this router as master (owner value 1).	`address vlan 1 2 10.0.0.2 1`
7.	Set a virtual router address as VLAN 1, VRID 1, interface 10.0.0.1, and set this router as backup (owner value 0).	`address vlan 1 1 10.0.0.1 0`

Procedure 9-4 Configuring Router 2 for Symmetrical VRRP (continued)

Step	Task	Command(s)
8.	Enable VRRP on this router on VLAN 1, VRIDs 1 and 2.	**enable vlan 1 1** **enable vlan 1 2**

Note: Before enabling VRRP on this router you must set the other options described in this section. Once enabled, you cannot make any configuration changes to VRRP on this router without first disabling VRRP using the **no enable vlan** command.

Example: Symmetrical VRRP Configuration

This example shows the complete configuration for symmetrical VRRP on Routers 1 and 2:

```
**Router 1**
Matrix>Router1(config)#interface vlan 1
Matrix>Router1(config-if(Vlan 1))#ip address 10.0.0.1 255.255.255.0
Matrix>Router1(config-if(Vlan 1))#no shutdown
Matrix>Router1(config)#router vrrp
Matrix>Router1(config-router)#create vlan 1 1
Matrix>Router1(config-router)#create vlan 1 2
Matrix>Router1(config-router)#address vlan 1 1 10.0.0.1 1
Matrix>Router1(config-router)#address vlan 1 2 10.0.0.2 0
Matrix>Router1(config-router)#enable vlan 1 1
Matrix>Router1(config-router)#enable vlan 1 2

**Router 2**
Matrix>Router2(config)#interface vlan 1
Matrix>Router2(config-if(Vlan 1))#ip address 10.0.0.2 255.255.255.0
Matrix>Router2(config-if(Vlan 1))#no shutdown
Matrix>Router2(config)#router vrrp
Matrix>Router2(config-router)#create vlan 1 1
Matrix>Router2(config-router)#create vlan 1 2
Matrix>Router2(config-router)#address vlan 1 2 10.0.0.2 1
Matrix>Router2(config-router)#address vlan 1 1 10.0.0.1 0
Matrix>Router2(config-router)#enable vlan 1 1
Matrix>Router2(config-router)#enable vlan 1 2
```

Multi-Backup Configuration

Figure 9-3 shows a VRRP configuration with three routers and three virtual routers. Each router serves as a master for one virtual router and as a backup for each of the others. When a master router goes down, one of the backups takes over the IP addresses of its virtual router.

In a VRRP configuration where more than one router is backing up a master, you can specify which backup router takes over when the master goes down by setting the priority for the backup routers.

Figure 9-3 Multi-Backup VRRP Configuration

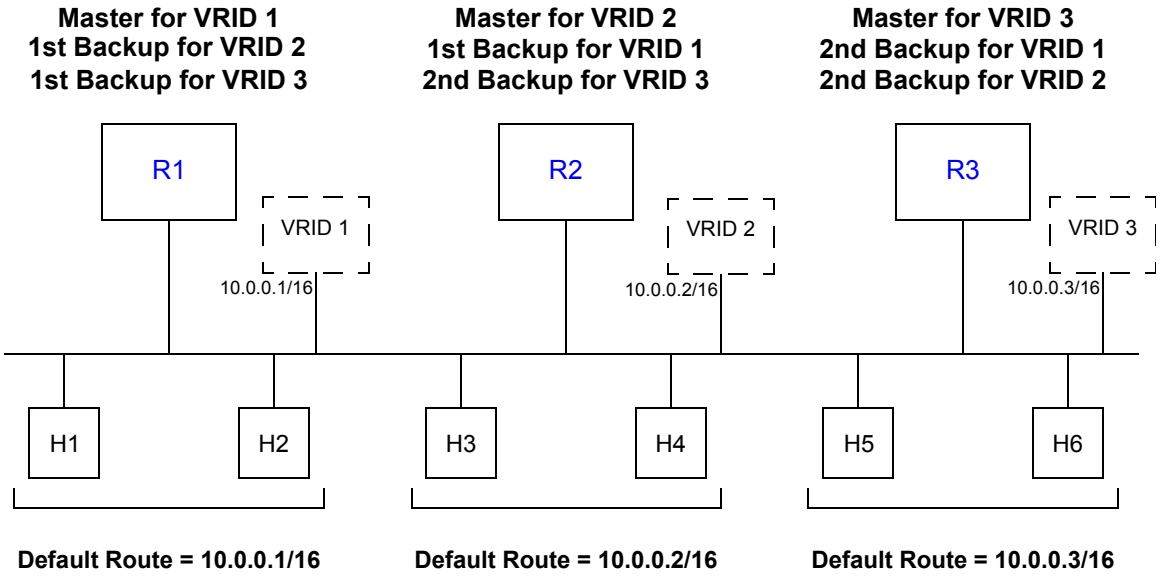

In this configuration, Router R1 is the master for virtual router VRID 1 and the primary backup for virtual routers VRID 2 and VRID 3. If Router R2 or R3 were to go down, Router R1 would assume the IP addresses associated with virtual routers VRID 2 and VRID 3.

Router R2 is the master for virtual router VRID 2, the primary backup for virtual router VRID 1, and the secondary backup for virtual router VRID 3. If Router R1 should fail, Router R2 would become the master for virtual router VRID 1. If both Routers R1 and R3 should fail, Router R2 would become the master for all three virtual routers. Packets sent to IP addresses 10.0.0.1/16, 10.0.0.2/16, and 10.0.0.3/16 would all go to Router R2.

Router R3 is the secondary backup for virtual routers VRID 1 and VRID 2. It would become a master router only if both Routers R1 and R2 should fail. In such a case, Router R3 would become the master for all three virtual routers.

Router R1 Configuration

Use the following procedure to configure Router R1 as shown back in Figure 9-3:

Procedure 9-5 Configuring Router 1 for Multi-Backup VRRP

Step	Task	Command(s)
1.	In **Global Configuration** mode: Create a routing interface for Router 1 on VLAN 1.	`interface vlan 1`
2.	In **Interface Configuration** mode: Set and enable IP address 10.0.0.1 255.255.255.0 on this router and VLAN.	`ip address 10.0.0.1 255.255.255.0` `no shutdown`

Procedure 9-5 Configuring Router 1 for Multi-Backup VRRP (continued)

Step	Task	Command(s)
3.	In **Global Configuration** mode: Enable VRRP configuration mode on this router.	`router vrrp`
4.	In **VRRP Configuration** Mode: Create a VRRP session for this router on VLAN 1 with a VRID of 1.	`create vlan 1 1`
5.	Create a second VRRP with a VRID of 2.	`create vlan 1 2`
6.	Create a third VRRP session with a VRID of 3.	`create vlan 1 3`
7.	Set a virtual router address as VLAN 1, VRID 1, interface 10.0.0.1, and set this router as master (owner value 1).	`address vlan 1 1 10.0.0.1 1`
8.	Set a virtual router address as VLAN 1, VRID 2, interface 10.0.0.2, and set this router as backup (owner value 0).	`address vlan 1 2 10.0.0.2 0`
9.	Set a virtual router address as VLAN 1, VRID 3, interface 10.0.0.3, and set this router as backup (owner value 0).	`address vlan 1 3 10.0.0.3 0`
10.	Set a priority of 200 for this router on VRIDs 2 and 3, making it the primary backup.	`priority vlan 1 2 200` `priority vlan 1 3 200`

11.	Enable VRRP on this router on VLAN 1, VRIDs 1, 2, and 3.	`enable vlan 1 1`
		`enable vlan 1 2`
		`enable vlan 1 3`

Note: Before enabling VRRP on this router you must set the other options described in this section. Once enabled, you cannot make any configuration changes to VRRP on this router without first disabling VRRP using the **no enable vlan** command.

Table 9-2 shows the priorities for each virtual router configured on Router R1.

Table 9-2 Priorities for Virtual Routers Configured on Router 1

Virtual Router	Configured Priority	Default Priority	Comments
VRID 1 on IP address 10.0.0.1/16	255	255	Since Router R1 is the owner of the IP address associated with VRID 1, the default priority of 255 (highest) does not need to be reconfigured to make this router the VRID 1 master.
VRID 2 on IP address 10.0.0.2/16	200	100	Changing Router R1's priority from 100 to 200 makes this virtual router first backup for VRIDs 2 and 3. If no other routers in the VRRP configuration have a higher priority, Router R1 will take over as master for VRIDs 2 and 3 should Router R2 or R3 go down.
VRID 3 on IP address 10.0.0.3/16	200	100	

Router R2 Configuration

Use the following procedure to configure Router R2 as shown back in Figure 9-3:

Procedure 9-6 Configuring Router 2 for Multi-Backup VRRP

Step	Task	Command(s)
1.	In **Global Configuration** mode: Create a routing interface for Router 2 on VLAN 1.	`interface vlan 1`
2.	In **Interface Configuration** mode: Set and enable IP address 10.0.0.2 255.255.255.0 on this router and VLAN.	`ip address 10.0.0.2 255.255.255.0` `no shutdown`
3.	In **Global Configuration** mode: Enable VRRP configuration mode on this router.	`router vrrp`
4.	In **VRRP Configuration** mode: Create a VRRP session for this router on VLAN 1 with a VRID of 1.	`create vlan 1 1`
5.	Create a second VRRP with a VRID of 2.	`create vlan 1 2`
6.	Create a third VRRP session with a VRID of 3.	`create vlan 1 3`
7.	Set a virtual router address as VLAN 1, VRID 2, interface 10.0.0.2, and set this router as master (owner value 1).	`address vlan 1 2 10.0.0.2 1`
8.	Set a virtual router address as VLAN 1, VRID 1, interface 10.0.0.1, and set this router as backup (owner value 0).	`address vlan 1 1 10.0.0.1 0`
9.	Set a virtual router address as VLAN 1, VRID 3, interface 10.0.0.3, and set this router as backup (owner value 0).	`address vlan 1 3 10.0.0.3 0`
10.	Set a priority of 200 for this router on VRID 1, making it the primary backup.	`priority vlan 1 1 200`
11.	Set a priority of 100 for this router on VRID 3, making it the secondary backup. **Note:** This command is shown for illustration purposes only since 100 is the default priority value and doesn't have to be set.	`priority vlan 1 3 100`

12.	Enable VRRP on this router on VLAN 1, VRIDs 1, 2, and 3.
	Note: Before enabling VRRP on this router you must set the other options described in this section. Once enabled, you cannot make any configuration changes to VRRP on this router without first disabling VRRP using the **no enable vlan** command.
		`enable vlan 1 1`
		`enable vlan 1 2`
		`enable vlan 1 3`

Table 9-3 shows the priorities for each virtual router configured on Router R2.

Table 9-3 Priorities for Virtual Routers Configured on Router 2

Virtual Router	Configured Priority	Default Priority	Comments
VRID 1 on IP address 10.0.0.1/16	200	100	Changing Router R2's priority from 100 to 200 makes this virtual router primary backup for VRID 1. Since this number is higher than Router R3's priority for VRID 1, R3 is the secondary backup.
VRID 2 on IP address 10.0.0.2/16	255	255	Since Router R2 is the owner of the IP address associated with VRID 2, the default priority of 255 (highest) does not need to be reconfigured to make this router the VRID 2 master.
VRID 3 on IP address 10.0.0.3/16	100	100	Leaving priority at 100 (lower than R1's) makes Router R2 secondary backup and R1 primary backup for VRID 3.

Router R3 Configuration

Use the following procedure to configure Router R3 as shown back in Figure 9-3:

Procedure 9-7 Configuring Router 3 for Multi-Backup VRRP

Step	Task	Command(s)
1.	In **Global Configuration** mode: Create a routing interface for Router 3 on VLAN 1.	`interface vlan 1`
2.	In **Interface Configuration** mode: Set and enable IP address 10.0.0.3 255.255.255.0 on this router an VLAN.	`ip address 10.0.0.3 255.255.255.0` `no shutdown`
3.	In **Global Configuration** mode: Enable VRRP configuration mode on this router.	`router vrrp`
4.	In **VRRP Configuration** mode: Create a VRRP session for this router on VLAN 1 with a VRID of 1.	`create vlan 1 1`
5.	Create a second VRRP session with a VRID of 2.	`create vlan 1 2`
6.	Create a third VRRP session with a VRID of 3.	`create vlan 1 3`
7.	Set a virtual router address as VLAN 1, VRID 3, interface10.0.0.3, and set this router as master (owner value 1).	`address vlan 1 3 10.0.0.3 1`
8.	Set a virtual router address as VLAN 1, VRID 1, interface10.0.0.1, and set this router as backup (owner value 0).	`address vlan 1 1 10.0.0.1 0`
9.	Set a virtual router address as VLAN 1, VRID 2, interface10.0.0.2, and set this router as backup (owner value 0).	`address vlan 1 2 10.0.0.2 0`

Procedure 9-7 Configuring Router 3 for Multi-Backup VRRP (continued)

Step	Task	Command(s)
10.	Set a priority of 100 for this router on VRIDs 1 and 2, making it a secondary backup. **Note:** This command is shown for illustration purposes only since 100 is the default priority value and doesn't have to be set.	**priority vlan 1 1 100** **priority vlan 1 2 100**
11.	Enable VRRP on this router on VLAN 1, VRIDs 1, 2, and 3. **Note:** Before enabling VRRP on this router you must set the other options described in this section. Once enabled, you cannot make any configuration changes to VRRP on this router without first disabling VRRP using the **no enable vlan** command.	**enable vlan 1 1** **enable vlan 1 2** **enable vlan 1 3**

Table 9-4 shows the priorities for each virtual router configured on Router R3.

Table 9-4 Priorities for Virtual Routers Configured on Router 3

Virtual Router	Configured Priority	Default Priority	Comments
VRID 1 on IP address 10.0.0.1/16	100	100	Leaving priority at 100 makes Router R3 secondary backup for VRIDs 1 and 2.
VRID 2 on IP address 10.0.0.2/16	100	100	
VRID 3 on IP address 10.0.0.3/16	255	255	Since Router R3 is the owner of the IP address associated with VRID 3, the default priority of 255 (highest) does not need to be reconfigured to make this router the VRID 3 master.

Example: Multi-Backup VRRP Configuration

This example shows the complete configuration for multi-backup VRRP on Routers 1, 2 and 3:

```
**Router 1**
Matrix>Router1(config)#interface vlan 1
Matrix>Router1(config-if(Vlan 1))#ip address 10.0.0.1 255.255.255.0
Matrix>Router1(config-if(Vlan 1))#no shutdown
Matrix>Router1(config)#router vrrp
Matrix>Router1(config-router)#create vlan 1 1
Matrix>Router1(config-router)#create vlan 1 2
Matrix>Router1(config-router)#create vlan 1 3
Matrix>Router1(config-router)#address vlan 1 1 10.0.0.1 1
Matrix>Router1(config-router)#address vlan 1 2 10.0.0.2 0
Matrix>Router1(config-router)#address vlan 1 3 10.0.0.3 0
Matrix>Router1(config-router)#priority vlan 1 2 200
Matrix>Router1(config-router)#priority vlan 1 3 200
Matrix>Router1(config-router)#enable vlan 1 1
Matrix>Router1(config-router)#enable vlan 1 2
Matrix>Router1(config-router)#enable vlan 1 3
```

```
**Router 2**
Matrix>Router2(config)#interface vlan 1
Matrix>Router2(config-if(Vlan 1))#ip address 10.0.0.2 255.255.255.0
Matrix>Router2(config-if(Vlan 1))#no shutdown
Matrix>Router2(config)#router vrrp
Matrix>Router2(config-router)#create vlan 1 1
Matrix>Router2(config-router)#create vlan 1 2
Matrix>Router2(config-router)#create vlan 1 3
Matrix>Router2(config-router)#address vlan 1 2 10.0.0.2 1
Matrix>Router2(config-router)#address vlan 1 1 10.0.0.1 0
Matrix>Router2(config-router)#address vlan 1 3 10.0.0.3 0
Matrix>Router2(config-router)#priority vlan 1 1 200
Matrix>Router2(config-router)#priority vlan 1 3 100
Matrix>Router2(config-router)#enable vlan 1 1
Matrix>Router2(config-router)#enable vlan 1 2
Matrix>Router2(config-router)#enable vlan 1 3

**Router 3**
Matrix>Router3(config)#interface vlan 1
Matrix>Router3(config-if(Vlan 1))#ip address 10.0.0.3 255.255.255.0
Matrix>Router3(config-if(Vlan 1))#no shutdown
Matrix>Router3(config)#router vrrp
Matrix>Router3(config-router)#create vlan 1 1
Matrix>Router3(config-router)#create vlan 1 2
Matrix>Router3(config-router)#create vlan 1 3
Matrix>Router3(config-router)#address vlan 1 3 10.0.0.3 1
Matrix>Router3(config-router)#address vlan 1 1 10.0.0.1 0
Matrix>Router3(config-router)#address vlan 1 2 10.0.0.2 0
Matrix>Router3(config-router)#priority vlan 1 1 100
Matrix>Router3(config-router)#priority vlan 1 2 100
Matrix>Router3(config-router)#enable vlan 1 1
Matrix>Router3(config-router)#enable vlan 1 2
Matrix>Router3(config-router)#enable vlan 1 3
```

Modifying a Configuration

This section describes optional settings and settings you can modify in a VRRP configuration, including backup priority, advertisement interval, critical IP address, pre-empt mode, and authentication key.

Setting the Backup Priority

As described in "Multi-Backup Configuration" on page 9-8, you can specify which backup router takes over when the master router goes down by setting the priority for the backup routers.

Use the following command in router configuration mode to set the priority for a backup router:

priority vlan *vlan-id vrid priority-value*

The priority value can be between 1 (lowest) and 254. The default is 100. The priority for the IP address owner is 255 and cannot be changed. Priority value of 255 is reserved for the VRRP router that owns the IP address associated with the virtual router. Priority 0 is reserved for signaling that the master has stopped working and the backup router must transition to master state.

Setting the Advertisement Interval

The VRRP master router sends periodic advertisement messages to let the other routers know that the master is up and running. By default, advertisement messages are sent once each second.

Use the following command in router configuration mode to change the VRRP advertisement interval:

```
advertise-interval vlan vlan-id vrid interval
```

Valid interval values are from 1 to 255 seconds.

Setting a Critical IP Address

In VRRP, a critical IP address defines an interface — in addition to the interface between hosts and a first-hop router — that will prevent the master router from functioning properly if the interface were to fail. For example, an IP address of an interface connecting a master router to a router configured for internet access would be considered a critical IP address for VRRP routing.

Use the following command in router configuration mode to set a critical IP address for VRRP routing:

```
critical-ip vlan vlan-id vrid ip-address
```

Setting Pre-empt Mode

When a master router goes down, the backup with the highest priority takes over the IP addresses associated with the master. By default, when the original master comes back up again, it takes over from the backup router that assumed its role as master. When a VRRP router does this, it is said to be in pre-empt mode. Pre-empt mode is enabled by default on the Matrix DFE Series or N-SA device. You can prevent a VRRP router from taking over from a lower-priority master by disabling pre-empt mode.

Use the following command in router configuration mode to disable pre-empt mode. This will prevent a backup router from taking over as master from a master router that has a lower priority:

```
no preempt vlan-id vrid
```

Note: If a VRRP configured IP address owner is available, it will always take over as the master router—regardless of whether pre-empt mode is on or off.

Setting an Authentication Key

VRRP Packet Authentication Fields

As shown in the following illustration, each VRRP packet contains authentication fields in its IP header. Table 9-5 describes VRRP packet authentication fields and their descriptions.

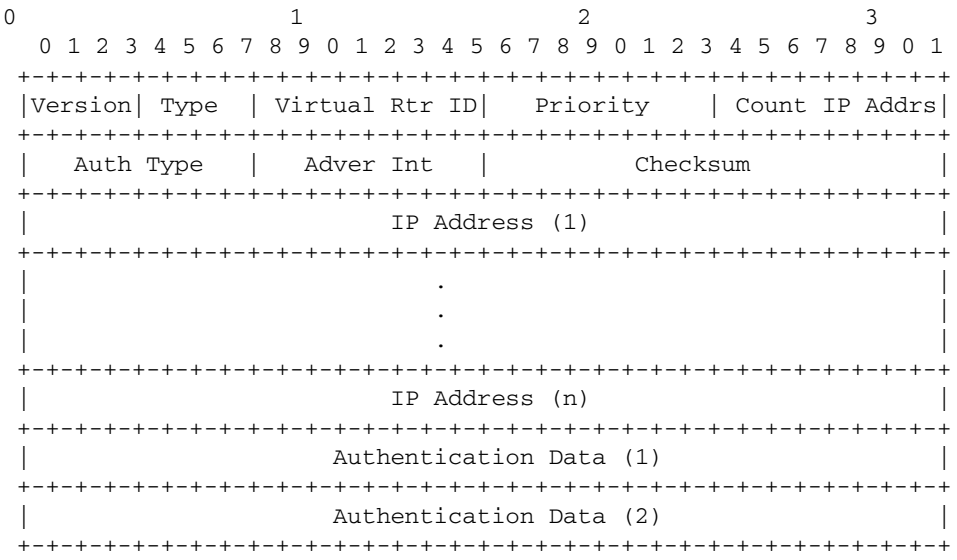
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
Version
Type
Virtual Rtr ID
Priority
Count IP Addrs
Auth Type
Adver Int
Checksum
IP Address (1)
IP Address (n)
Authentication Data (1)
Authentication Data (2)

Table 9-5 VRRP Packet Authentication Field Descriptions

Field	Description
`Auth Type`	Identifies the authentication method being utilized. Authentication type is an 8-bit unsigned integer and is unique on a per interface basis. A VRRP packet with an unknown authentication type, or that does not match the locally configured authentication method, must be discarded. Authentication methods currently defined are: • Authentication Data Type 0 - No authentication (default setting) • Authentication Data Type 1 - Simple Text Password • Authentication Data Type 2 - Message Digest (MD5) Corresponding contents of the Authentication Data fields for each Authentication Type are listed below.
`Authentication Data (1)`	Simple Text Password. The contents of the Authentication Data field should be set to the locally configured password on transmission. There is no default password. The receiver must check that the Authentication Data in the packet matches its configured authentication string. Packets that do not match must be discarded. This password is a text string 1 to 8 characters in length and can be set with the i**p vrrp authentication-key** command.
`Authentication Data (2)`	Message Digest (MD5) key. If a packet is received that does not pass the authentication check due to a missing authentication header or incorrect message digest, then the packet must be discarded. This password is a text string 1 to 16 characters in length and can be set with the i**p vrrp message-digest-key** command.

VRRP Authentication Commands

By default, no authentication of VRRP packets is performed on the Matrix DFE Series or N-SA device. You can specify a clear text or a message digest (MD5) password to be used to authenticate VRRP exchanges on routing interfaces.

Use the following command in interface configuration mode to set a clear text authentication password:

ip vrrp authentication-key password

Password is a text string 1 to 8 characters in length.

Use the following command in interface configuration mode to set an MD5 authentication password:

ip vrrp message-digest-key *vrid* **md5** *password*

Password is a text string 1 to 16 characters in length.

Monitoring VRRP

Use the following command in global configuration mode to display VRRP routing information:

show ip vrrp

Example

This example shows that Router 1 is configured on VLAN 1, VRID 1, as the master router of VRID 1 and the owner of IP address 10.0.0.1. It also shows Router 1 is configured as the primary backup router for VRIDs 2 and 3:

```
Matrix>Router1(config)#show ip vrrp
-----------VRRP CONFIGURATION-----------
Vlan    Vrid       State          Owner     AssocIpAddr
1       1          Master           1       10.0.0.1
1       2          Primary          0       10.0.0.2
1       3          Primary           0       10.0.0.3
```

VRRP Configuration Notes

Determining the Backup to Master Time Interval

If a backup router doesn't receive a keep-alive advertisement from the current master within a certain period of time, it will transition to the master state and start sending advertisements to itself. The amount of time that a backup router will wait before it becomes the new master is based on the following equation:

Master-down-interval = (3 * advertisement-interval) + skew-time

The skew-time depends on the backup router's configured priority:

Skew-time = ((256 - Priority) / 256)

Therefore, the higher the priority, the faster a backup router will detect that the master is down. For example:

- Default advertisement-interval = 1 second

- Default backup router priority = 100
- Master-down-interval = time it takes a backup to detect the master is down
  - = (3 * adv-interval) + skew-time
  - = (3 * 1 second) + ((256 - 100) / 256)
  - = 3.6 seconds

Note: In some instances, a heavy load on a VRRP master may delay VRRP packet transmission and cause the backup router to assume the role of master.

When the Master Reboots or Is Down

If a master router is manually rebooted, or if its interface is manually brought down, it will send a special keep-alive advertisement that lets the backup routers know that a new master is needed immediately.

Determining a Virtual MAC Address

A virtual router will respond to ARP requests with a virtual MAC address. This virtual MAC depends on the virtual router ID:

virtual MAC address = `00005E:0001XX`

where *XX* is the virtual router ID

This virtual MAC address is also used as the source MAC address of the keep-alive Advertisements transmitted by the master router.

- These MAC addresses, when active, use entries in the port's Routing Address Table (RAT). Since the RAT has eight entries and other resources also use the RAT, the system will limit the number of unique virtual routing ID's to six or fewer on a per port basis.
  - If multiple virtual routers are created on a single interface, the virtual routers must have unique identifiers. If virtual routers are created on different interfaces, you can reuse virtual router IDs.
  - As specified in RFC 2338, a backup router that has transitioned to master will not respond to pings, accept telnet sessions, or field SNMP requests directed at the virtual router's IP address. Not responding allows network management to notice that the original master router (that is, the IP address owner) is down.